Robust information propagation through noisy neural circuits

Joel Zylberberg[1,2,3,4,*], Alexandre Pouget[5,6], Peter E. Latham[6,†], Eric Shea-Brown[3,7,8,†]

1 Department of Physiology and Biophysics, Center for Neuroscience, and Computational Bioscience Program, University of Colorado School of Medicine, Aurora, Colorado, United States of America
2 Department of Applied Mathematics, University of Colorado, Boulder, Colorado, United States of America
3 Department of Applied Mathematics, University of Washington, Seattle, Washington, United States of America
4 Learning in Machines and Brains Program, Canadian Institute For Advanced Research, Toronto, Ontario, Canada
5 Department of Basic Neuroscience, University of Geneva, Switzerland
6 Gatsby Computational Neuroscience Unit, University College London, London, United Kingdom
7 Department of Physiology and Biophysics, Program in Neuroscience, University of Washington Institute for Neuroengineering, and Center for Sensorimotor Neural Engineering, University of Washington, Seattle, Washington, United States of America
8 Allen Institute for Brain Science, Seattle, Washington, United States of America
† These authors contributed equally to this work
∗ Email: joel.zylberberg@ucdenver.edu

arXiv:1608.05706v2 [q-bio.NC] 26 Feb 2017

Abstract

Sensory neurons give highly variable responses to stimulation, which can limit the amount of stimulus information available to downstream circuits. Much work has investigated the factors that affect the amount of information encoded in these population responses, leading to insights about the role of covariability among neurons, tuning curve shape, etc. However, the informativeness of neural responses is not the only relevant feature of population codes; of potentially equal importance is how robustly that information propagates to downstream structures. For instance, to quantify the retina's performance, one must consider not only the informativeness of the optic nerve responses, but also the amount of information that survives the spike-generating nonlinearity and noise corruption in the next stage of processing, the lateral geniculate nucleus. Our study identifies the set of covariance structures for the upstream cells that optimize the ability of information to propagate through noisy, nonlinear circuits. Within this optimal family are covariances with "differential correlations", which are known to reduce the information encoded in neural population activities. Thus, covariance structures that maximize information in neural population codes, and those that maximize the ability of this information to propagate, can be very different. Moreover, redundancy is neither necessary nor sufficient to make population codes robust against corruption by noise: redundant codes can be very fragile, and synergistic codes can – in some cases – optimize robustness against noise.

Author Summary

Information about the outside world, which originates in sensory neurons, propagates through multiple stages of processing before reaching the neural structures that control behavior. While much work in neuroscience has investigated the factors that affect the amount of information contained in peripheral sensory areas, very little work has asked how much of that information makes it through subsequent processing stages. That's the focus of this paper, and it's an important issue because information that fails to propagate cannot be used to affect decision-making. We find a tradeoff between information content and information transmission: neural codes which contain a large amount of information can transmit that information poorly to subsequent processing stages. Thus, the problem of robust information propagation

– which has largely been overlooked in previous research – may be critical for determining how our sensory organs communicate with our brains. We identify the conditions under which information propagates well – or poorly – through multiple stages of neural processing.

1 Introduction

Neurons in sensory systems gather information about the environment, and transmit that information to other parts of the nervous system. This information is encoded in the activity of neural populations, and that activity is variable: repeated presentations of the same stimulus lead to different neuronal responses [1–7]. This variability can degrade the ability of neural populations to encode information about stimuli, leading to the question: which features of population codes help to combat – or exacerbate – information loss?

This question is typically addressed by assessing the amount of information that is encoded in the periphery as a function of the covariance structure [6, 8–24], the shapes of the tuning curves [25, 26], or both [27, 28]. However, the informativeness of the population responses at the periphery is not the only relevant quantity for understanding sensory coding; of potentially equal importance is the amount of information that propagates through the neural circuit to downstream structures [29, 30].

To illustrate the ideas, consider the case of retinal ganglion cells transmitting information about visual stimuli to the cortex via the thalamus, as shown in Fig. 1. To quantify the performance of the retina, one must consider not only the informativeness of the optic nerve responses ($I_x(s)$ in Fig. 1A), but also how much of that information is transmitted by the lateral geniculate nucleus (LGN) to the cortex ($I_y(s)$ in Fig. 1A) [31]. The two may be very different, as only information that survives the LGN's spike-generating nonlinearity and noise corruption will propagate to downstream cortical structures.

Despite its importance, the ability of information to propagate through neural circuits remains relatively unexplored [31]. One notable exception is the literature on how synchrony among the spikes of different cells affects responses in downstream populations [32–36]. This is, however, distinct from the information propagation question we consider here, as there is no guarantee that those downstream spikes will be informative. Other work [25, 29, 30, 37, 38] investigated the question of optimal network properties (tuning curves and connection matrices) for information propagation in the presence of noise.

No prior work, however, has isolated the impact of correlations on the ability of population-coded information to propagate. Given the frequent observations of correlations in the sensory periphery [6, 8, 17, 39–45], and the importance of the information propagation problem, this is a significant gap in our knowledge. To fill that gap, we consider a model (Fig. 1B; described in more detail below), in which there are two layers (retina and LGN, for example). The first layer contains a fixed amount of information, $I_x(s)$, which is encoded in the noisy, stimulus-dependent responses of the cells in that layer. The information is passed to the second layer via feedforward connections followed by a nonlinearity, with noise added along the way. We ask how the covariance structure of the trial-to-trial variability in the first layer affects the amount of information in the second.

Although we focus on information propagation, the problem we consider applies to more general scenarios. In essence, we are asking: how does the noise in the input to a network interact with noise added to the output? Because we consider linear feedforward weights followed by a nonlinearity, the possible transformations from input to output, and thus the computations the network could perform, is quite broad [46]. Thus, the conclusions we draw apply not just to information propagation, but also to many computations. Moreover, it may be possible to extend our analysis to recurrent, time-dependent neural networks. That is, however, beyond the scope of this work.

Our results indicate that the amount of information that successfully propagates to the second layer depends strongly on the structure of correlated responses in the first. For linear neural gain functions, and some classes of nonlinear ones, we identify analytically the covariance structures that optimize information propagation through noisy downstream circuits. Within the optimal family of covariance structures, we

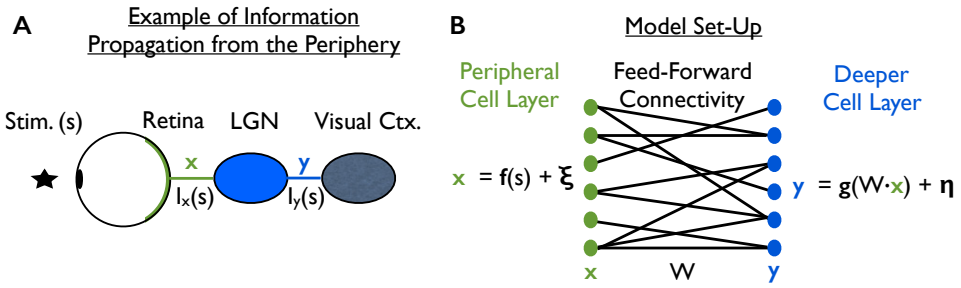

Figure 1. The information propagation problem. This problem is illustrated with the visual periphery, but the information propagation problem is general: it arises whenever information is transmitted from one area to another, and also when information is combined to carry out computations. (**A**) The retina transmits information about visual stimuli, s, to the visual cortex. The information does not propagate directly from retina to cortex; it is transmitted via an intermediary structure, the lateral geniculate nucleus (LGN). Consequently, the information about the stimuli that is available to the cortex, denoted $I_y(s)$, is not the same as the information that retina transmits, denoted $I_x(s)$. Here, we ask what properties of neural activities in the periphery maximize the information that propagates to the deeper neural structures. (**B**) Illustration of our model. Neural activity in the periphery, $\mathbf{x}$, is generated by passing the stimulus, s, through a set of neural tuning curves, $\mathbf{f}(s)$, and then adding zero-mean noise, $\boldsymbol{\xi}$, which may be correlated between cells. This activity then propagates via feed-forward connectivity, described by the matrix $\mathbf{W}$, to the next layer. The activity at the next layer, $\mathbf{y}$, is generated by passing the inputs, $\mathbf{W} \cdot \mathbf{x}$, through a nonlinearity $g(\cdot)$, and then adding zero-mean noise, $\boldsymbol{\eta}$.

find variability with so-called differential correlations [22] – correlations that are proven to minimize the information in neural population activity. Thus, covariance structures that maximize the information content of neural population codes, and those that maximize the ability of this information to propagate, can be very different. Importantly, we also find that redundancy is neither necessary nor sufficient for the population code to be robust against corruption by noise. Consequently, to understand how correlated neural activity affects the function of neural systems, we must not only consider the impact of those correlations on information, but also the ability of the encoded information to propagate robustly through multi-layer circuits.

2 Results

2.1 Problem Formulation: Information Propagation in the Presence of Corrupting Noise

We consider a model in which a vector of "peripheral" neural population responses, $\mathbf{x}$, is determined by two components. The first is the set of tuning curves, $\mathbf{f}(s)$, which define the cells' mean responses to any particular stimulus (typical tuning curves are shown in Fig. 2A). Here we consider a one dimensional stimulus, denoted s, which may represent, for example, the direction of motion of a visual object. In that case, a natural interpretation of our model is that it describes the transmission of motion information by direction selective retinal ganglion cells to the visual cortex (Fig. 1) [5, 6, 47]. Extension to multi-dimensional stimuli is straightforward. The second component of the neural population responses, $\boldsymbol{\xi}$, represents the trial-to-trial variability. This results in the usual "tuning curve plus noise" model,

$$\mathbf{x} = \mathbf{f}(s) + \boldsymbol{\xi}, \tag{1}$$

where $\boldsymbol{\xi}$ is a zero mean random variable with covariance $\boldsymbol{\Sigma}_\xi$.

The neural activity, $\mathbf{x}$, propagates to the second layer via feed-forward weights, $\mathbf{W}$, as in the model of [38]. The activity in the second layer is given by passing the input, $\mathbf{W} \cdot \mathbf{x}$, through a nonlinearity, $g(\cdot)$, and then corrupting it with noise, $\boldsymbol{\eta}$ (Fig. 1B),

$$\mathbf{y} = g(\mathbf{W} \cdot \mathbf{x}) + \boldsymbol{\eta}, \tag{2}$$

where the nonlinearity is taken component by component, and $\boldsymbol{\eta}$ is zero mean noise with covariance matrix $\boldsymbol{\Sigma}_\eta$. The function $g(\cdot)$ need not be invertible, so this model can include spike generation.

While we have, in Fig. 1, given one explicit interpretation of our model, the model itself is quite general. This means that our results apply more broadly than just to circuits in the peripheral visual system. Moreover, while our analysis (below) focuses on information loss between layers, this should not be taken to mean that there is no meaningful computation happening within the circuit: because we have considered arbitrary nonlinear transformations between layers, the same model can describe a wide range of possible computations [46]. Our results apply to information loss during those computations.

In the standard fashion [6, 12, 20–22], we quantify the information in the neural responses using the linear Fisher information. This measure quantifies the precision (inverse of the mean squared error) with which a locally optimal linear estimator can recover the stimulus from the neural responses [48, 49]. The linear Fisher information in the first and second layers, denoted $I_x(s)$ and $I_y(s)$, respectively, is given by

$$I_x(s) = \mathbf{f}'(s) \cdot \boldsymbol{\Sigma}_\xi^{-1} \cdot \mathbf{f}'(s) \tag{3a}$$

$$I_y(s) = \mathbf{f}'(s) \cdot \left[\boldsymbol{\Sigma}_\xi + (\mathbf{W}_{\text{eff}}^T \cdot \boldsymbol{\Sigma}_{\text{eff},\eta}^{-1} \cdot \mathbf{W}_{\text{eff}})^{-1}\right]^{-1} \cdot \mathbf{f}'(s) \tag{3b}$$

where a prime denotes a derivative. Here $\mathbf{W}_{\text{eff}}$ are the effective weights – basically, the weights, $\mathbf{W}$, multiplied by the average slope of the gain function, $g(\cdot)$ – and $\boldsymbol{\Sigma}_{\text{eff},\eta}$ includes contributions from the noise in the second layer, $\boldsymbol{\eta}$, and, if $g(\cdot)$ is nonlinear, from the noise in the first layer. (If g is linear, $\boldsymbol{\Sigma}_{\text{eff},\eta} = \boldsymbol{\Sigma}_\eta$, so in this case $\boldsymbol{\Sigma}_{\text{eff},\eta}$ depends only on the noise in the second layer). This expression is valid if $\mathbf{W}_{\text{eff}}^T \cdot \boldsymbol{\Sigma}_{\text{eff},\eta}^{-1} \cdot \mathbf{W}_{\text{eff}}$ is invertible; so long as there are more cells in the second layer than the first, this is typically the case. See Methods, Sec. 4.1, for details.

Equation (3b) is somewhat intuitive, at least at a gross level: both large effective noise ($\boldsymbol{\Sigma}_{\text{eff},\eta}$) and small effective weights ($\mathbf{W}_{\text{eff}}$) reduce the amount of information at the second layer. At a finer level, the relationship between the two covariance structures – corresponding to the first and second terms in brackets in Eq. (3b) – can have a large effect on $I_y(s)$, as we will see shortly.

2.2 Information content and information propagation put different constraints on neural population codes

We begin with an example to highlight the difference between the information contained in neural population codes and the information that propagates through subsequent layers. Here, we consider two different neuronal populations with identical tuning curves (Fig. 2A), nearly-identical levels of trial-to-trial neural variability, and identical amounts of stimulus information encoded in their firing-rate responses; the populations' correlational structures, however, differ. We then corrupt these two populations' response patterns with noise, to mimic corruption that might arise in subsequent processing stages, and ask how much of the stimulus information remains. Surprisingly, the two population codes can show very different amounts of information after corruption by even modest amounts of noise (Fig. 2B).

In more detail, there are 100 neurons in the first layer; those neurons encode an angle, denoted s, via their randomly-shaped and located tuning curves (Fig. 2A). We consider two separate model populations.

Both have the same tuning curves, but different covariance matrices. For reasons we discuss below, those covariance matrices, denoted $\mathbf{\Sigma}_\xi^{\text{blue}}$ and $\mathbf{\Sigma}_\xi^{\text{green}}$ (blue and green correspond to the colors in Figs. 2B and C), are given by

$$\mathbf{\Sigma}_\xi^{\text{blue}} = \mathbf{\Sigma}_0 + \epsilon \mathbf{f}'(\mathbf{s})\mathbf{f}'(\mathbf{s}) \tag{4a}$$

$$\mathbf{\Sigma}_\xi^{\text{green}} = \mathbf{\Sigma}_0 + \epsilon_u \mathbf{u}(s)\mathbf{u}(s) \tag{4b}$$

where $\mathbf{\Sigma}_0$ is a diagonal matrix with elements equal to the mean response,

$$\Sigma_{0,ij} = f_i(s)\delta_{ij}. \tag{5}$$

Here δ_{ij} is the Kronecker delta ($\delta_{ij} = 1$ if $i = j$ and 0 otherwise), and we use the convention that two adjacent vectors denote an outer product; for instance, the ij^{th} element if $\mathbf{uu}$ is $u_i u_j$. The vector $\mathbf{u}$ has the same magnitude as $\mathbf{f}'$, but points in a slightly different direction (it makes an angle θ_u with $\mathbf{f}'$), and ϵ and ϵ_u are chosen so that the information in the two populations, $I_x(s)$, is the same (ϵ_u also depends on s; we suppress that dependence for clarity).

In our simulations, both ϵ and ϵ_u are small (on the order of 10^{-3}; see Sec. 4.8.1), so the variance of the i^{th} neuron is approximately equal to its mean. This makes the variability Poisson-like, as is typically observed when counting neural spikes in finite time windows [1–6]. (More precisely, the average Fano factors – averaged over neurons and stimuli – were 1.01 for the "blue" population and 1.04 for the "green" one.) Both model populations also have the same average correlation coefficients, which are near-zero (see Methods, Sec. 4.8.1, for details).

To determine how much of the information in the two populations propagates to the second layer, we computed $I_y(s)$ for both populations using Eq. (3b). For simplicity, we used the identity matrix for the feed-forward weights, $\mathbf{W}$, a linear gain function, $g(\cdot)$, and independently and identically distributed (*iid*) noise with variance σ^2. Later we consider the more general case: arbitrary feedforward weights, nonlinear gain functions, and arbitrary covariance for the second layer noise. Those complications don't, however, change the basic story.

Figure 2B shows the information in the output layer versus the level of output noise, σ^2, for the two populations. Blue and green curves correspond to the different covariance structures. Although the two populations have identical tuning curves, nearly-identical levels of trial-to-trial neural variability, and contain identical amounts of information about the stimulus, they differ markedly in the robustness of that information to corruption by noise in the second layer. Thus, quantifying the information content of neural population codes is not sufficient to characterize them: recordings from the first-layer cells of the two example populations in Fig. 2 would yield identical information about the stimulus, but the blue population has a greater ability to propagate that information downstream.

One possible explanation for the difference in robustness is that the information in the green population relies heavily on correlations, which are destroyed by a small amount of noise. To check this, we compared the information of the correlated neural populations to the information that would be obtained with the same tuning curves and levels of single neuron trial-to-trial variability, but no inter-neuronal correlations [11, 50, 51] (Fig. 2C). We find that removing the correlations actually *increases* the information in both populations (Fig. 2C; "Trial-Shuffled"), and by about the same amount, so this possible explanation cannot account for the difference in robustness. We also considered the case where the correlated responses carry more information than would be obtained from independent cells. We again found (similar to Fig. 2) that there could be substantial differences in the amount of information propagated by equally informative population codes (see Methods, Fig. 7).

These examples illustrate that merely knowing the amount of information in a population, or how that information depends on correlations in neural responses, doesn't tell us how much of that information will propagate to the next layer. In the remainder of this paper, we provide a theoretical explanation

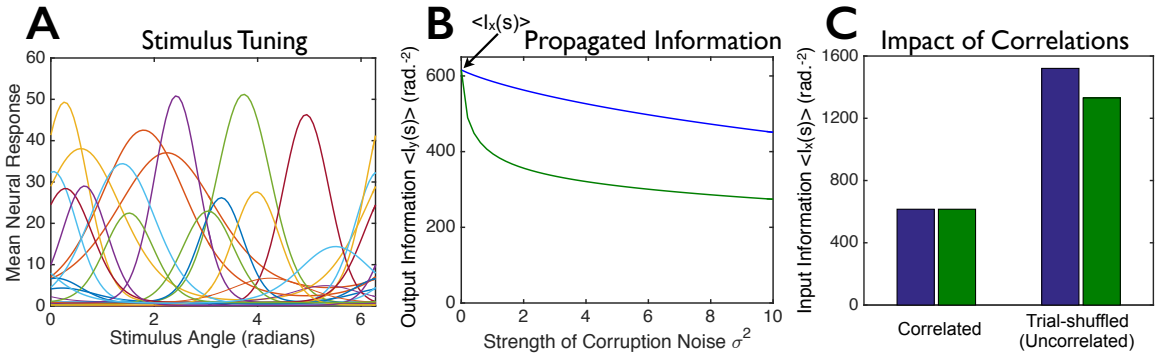

Figure 2. Not all population codes are equally robust against corruption by noise. We constructed two model populations, each with the same 100 tuning curves for the first layer of cells but with different covariance structures, $\mathbf{\Sigma}_\xi$ (see text, especially Eq. (4)). The covariance structures were chosen so that the two populations convey identical amounts of information $I_x(s)$ about the stimulus. (**A**) 20 randomly-chosen tuning curves from the 100 cell population. (**B**) We corrupted the responses of each neural population by additional Gaussian noise (independently and identically distributed for all cells) of variance σ^2, to mimic corruption that might arise as the signals propagate through a multi-layered neural circuit, and computed the "output" information $I_y(s)$ that these further-corrupted responses convey about the stimulus (blue and green curves). The population shown in green forms a relatively fragile code wherein modest amounts of noise strongly reduce the information, whereas the population shown in blue is more robust. (**C**) Input information $I_x(s)$ in the two model populations (left; "correlated") and information that would be conveyed by the model populations if they had their same tuning curves and levels of trial-to-trial variability, but no correlations between cells (right; "trial-shuffled"). For panels B and C, we computed the information for each of 100 equally spaced stimulus values, and averaged the information over those stimuli. See Methods, Sec. 4.8.1 for additional details.

of this observation, and identify the covariance structures at the first layer that maximize robustness to information loss during propagation through downstream circuits.

2.3 Geometry of Robust Versus Fragile Population Codes

To understand, from a geometrical point of view, why some population codes are more sensitive to noise than others, we need to consider the relationship between the noise covariance ellipse and the "signal direction," $\mathbf{f}'(s)$ – the direction the mean neural response changes when the stimulus s changes by a small amount. Figures 3A and B show this relationship for two different populations. The noise distribution in the first layer is indicated by the magenta ellipses, and the signal direction by the green arrows. The uncertainty in the stimulus after observing the neural response is indicated by the overlap of the green line with the magenta ellipse. Because the overlap is the same for the two populations, they have the same amount of stimulus uncertainty, and thus the same amount of information – at least in the first layer.

Although the two populations have the same amount of information, the covariance ellipses are very different: one long and skinny but slightly tilted relative to the signal direction (Fig. 3A), the other shorter and fatter and parallel to the signal direction (Fig. 3B). Consequently, when *iid* noise is added, as indicated by the dashed lines, stimulus uncertainty increases by very different amounts: there's a much larger increase for the long skinny ellipse than for the short fat one. This makes the population code in

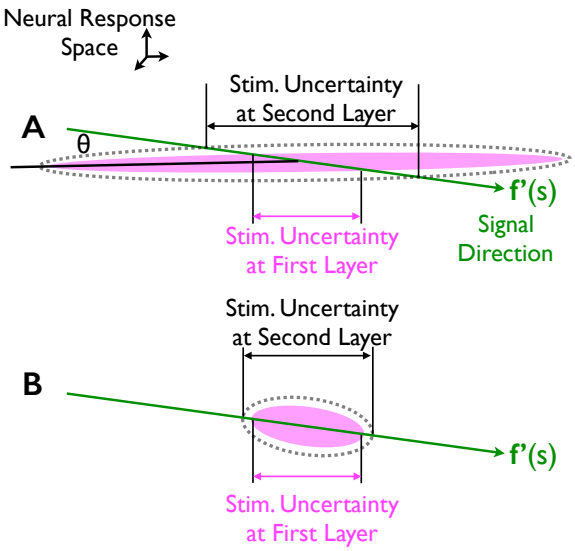

Figure 3. Geometry of robust versus fragile population codes. Cartoons showing the interaction of signal and noise for two populations with the same information in the input layer. The dimension of the space is equal to the number of cells in the population; we show a two dimensional projection. Within this space, when the stimulus changes by an amount Δs (with Δs small), the average neural response changes by $\mathbf{f}'(s)\Delta s$. Thus, $\mathbf{f}'(s)$ is the "signal direction" (green arrows). Trial-by-trial fluctuations in the neural responses in the first layer are described by the ellipses; these correspond to 1 standard-deviation probability contours of the conditional response distributions. The impact of the neural variability on the encoding of stimulus s is determined by the projection of the response distributions onto the signal direction (magenta double-headed arrows). By construction, these are identical in the first layer. Accordingly, an observer of the neural activity in the first layer of either population would have the same level of uncertainty about the stimulus, and so both populations encode the same amount of stimulus information. When additional *iid* noise is added to the neural responses, the response distributions grow; the dashed ellipses show the resultant response distributions at the second layer. Even though the same amount of *iid* noise is added to both populations, the one in panel A shows greater stimulus uncertainty after the addition of noise than does the one in panel B. Consequently, the information encoded by the population in panel B is more robust against corruption by noise.

Fig. 3A much more sensitive to added noise than the one in Fig. 3B.

To more rigorously support this intuition, in Methods, Sec. 4.4, we derive explicit expressions for the stimulus uncertainty in the first and second layers as a function of the angle between the long axis of the covariance ellipse and the signal direction. Those expressions corroborate the phenomenon shown in Fig. 3.

2.4 A family of optimal noise structures

The geometrical picture in the previous section tells us that a code is robust against added noise if the covariance ellipse lines up with the signal direction.Taken to it's extreme, this suggests that when all the noise is concentrated along the $\mathbf{f}'(s)$ direction, so that the covariance matrix is given by

$$\mathbf{\Sigma}_\xi(s) \propto \mathbf{f}'(s)\mathbf{f}'(s), \tag{6}$$

the resulting code should be optimally robust. While this may be intuitively appealing, the arguments that led to it were based on several assumptions: *iid* noise added in the second layer, feedforward weights, $\mathbf{W}$, set to the identity matrix, and a linear neural response function $g(\cdot)$. In real neural circuits, none of these assumptions hold. It turns out, though, that the only one that matters is the linearity of $g(\cdot)$. In this section we demonstrate that the covariance matrix given by Eq. (6) optimizes information transmission for neurons with linear gain functions (although we find, perhaps surprisingly, that this optimum is not unique). In the next section we consider nonlinear gain functions; for that case the covariance matrix given by Eq. (6) can be, but is not always guaranteed to be, optimal.

To determine what covariance structures maximize information propagation, we simply maximize information in the second layer, $I_y(s)$, with respect to the noise covariance matrix in the first layer, $\mathbf{\Sigma}_\xi$, with the information in the first layer held fixed. When the gain function, $g(\cdot)$, is linear (the focus of this section), this is relatively straightforward. Details of the calculation are given in Methods, Sec. 4.2; here we summarize the results.

The main finding is that there exists a family of first-layer covariance matrices $\mathbf{\Sigma}_\xi$, not just one, that maximizes the information in the second layer. That family, parameterized by α, is given by

$$\mathbf{\Sigma}_\xi(s) = \frac{\alpha}{I_x(s)}\, I_\eta(s)\mathbf{\Sigma}_y + \frac{1-\alpha}{I_x(s)}\, \mathbf{f}'(s)\mathbf{f}'(s), \tag{7}$$

where $\mathbf{\Sigma}_y$ is the effective covariance matrix in the second layer,

$$\mathbf{\Sigma}_y \equiv (\mathbf{W}_{\text{eff}}^T \cdot \mathbf{\Sigma}_\eta^{-1} \cdot \mathbf{W}_{\text{eff}})^{-1}, \tag{8}$$

and $I_\eta(s)$ is the information the second layer would have if there were no noise in the first layer,

$$I_\eta(s) = \mathbf{f}'(s) \cdot \mathbf{\Sigma}_y^{-1} \cdot \mathbf{f}'(s) \tag{9}$$

(see in particular Methods, Eq. (46)). For this whole family of distributions – that is, for any value of α for which $\mathbf{\Sigma}_\xi$ is positive semi-definite – the output information, $I_y(s)$, has exactly the same value,

$$I_y(s) = \frac{I_x(s)}{1 + I_x(s)/I_\eta(s)} \tag{10}$$

(see Methods, Eq. (76)). This is the maximum possible output information given the input information, $I_x(s)$.

Two members of this family are of particular interest. One is $\alpha = 0$, for which the covariance matrix corresponds to differential correlations (Eq. (6)); that covariance matrix is illustrated in Fig. 4A. This covariance matrix aligns the noise direction with the signal direction. Accordingly, as for the geometrical picture in Fig. 3, it makes the encoded information maximally robust.

The other family member we highlight is $\alpha = 1$, for which $\mathbf{\Sigma}_\xi \propto \mathbf{\Sigma}_y$. For this case, the covariance matrix in the first layer matches the effective covariance matrix in the second layer; we thus refer to this as "matched covariance". To understand why this covariance optimizes information in the second layer, we start with the observation that the population activities can be decomposed into their principal components: each principal component corresponds to a different axis along with the population activities can be projected. The information contained in each such projection (principal component) adds up to give the total Fisher information (see Methods, Eq. 71). The most informative of these projections are those that have low noise variance, and which align somewhat with the signal curve – like the blue line in Fig. 4B. When $\mathbf{\Sigma}_\xi \propto \mathbf{\Sigma}_y$, the projections that are most informative in the first layer are corrupted by relatively little noise in the second layer. Consequently, this configuration enables robust information propagation. In contrast, when the covariance structures in the first and second layers are less well

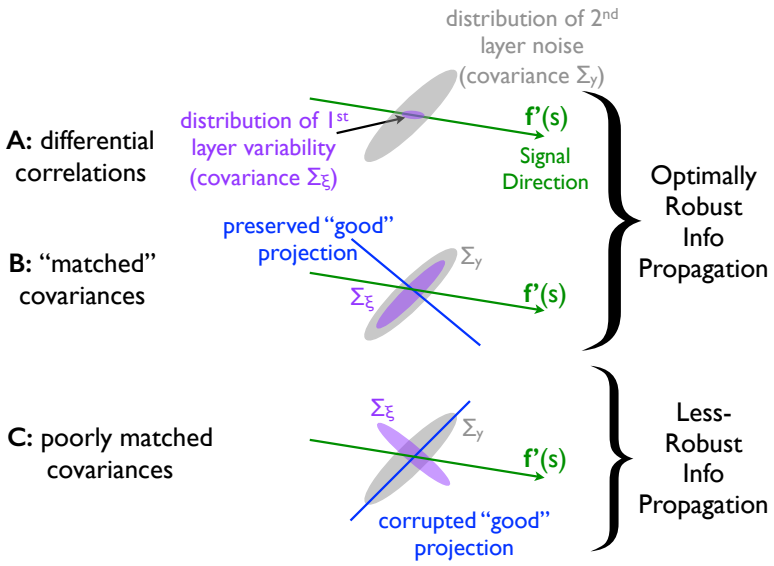

Figure 4. Family of optimal covariance matrices. For all panels, green arrows indicate the signal direction, $\mathbf{f}'(s)$. Magenta ellipses indicate the noise in the first layer (with corresponding covariance matrix $\mathbf{\Sigma}_\xi$), and grey ellipses indicate the effective noise in the second layer (with corresponding covariance matrix $\mathbf{\Sigma}_y$). (**A**) The covariance ellipse in the first layer has its long axis aligned with the signal direction; this configuration (which corresponds to differential correlations) optimizes information robustness for any distribution of second layer noise. (**B**) The covariance ellipse in the first layer does not have its long axis aligned with the signal direction. However, the covariance ellipse of the effective noise in the second layer, $\mathbf{\Sigma}_y$, has the same shape as the covariance ellipse in the first. In this case, the blue "good" projection – which is aligned both with a low-variance direction of the first-layer distribution (magenta), and with the signal curve (green), and thus is relatively informative about the stimulus (see text) – is corrupted by relatively little noise at the second layer. This "matched" noise configuration is among those that optimize robustness to noise. The optimal family of covariance matrices interpolates between the configurations shown in panels A and B. (**C**) Again the covariance ellipse in the first layer does not have its long axis aligned with the signal direction. But now the "good" projection is heavily corrupted by noise at the second layer. In this configuration, all projections are substantially corrupted by noise at some point in the circuit, and thus relatively little information can propagate.

matched, all projections are heavily corrupted by noise at some point (i.e., either in the first or the second layer), and hence very little information propagates (Fig. 4C).

The family of optima interpolates between the two configurations shown in Figs. 4A and B (see also Eq. (7)). Almost all members of this optimal covariance family depend on the details of the downstream circuit: for $\alpha \neq 0$ in Eq. (7), the optimal noise covariance at the first layer depends on the feed-forward weights, $\mathbf{W}$, and the structure of the downstream noise. The one exception to this is the covariance matrix given by Eq. (6): that one is optimal regardless of the downstream circuit. These are so-called "differential correlations" – the only correlations that lead to information saturation in large populations [22], and the correlations that minimize information in general (Methods, Sec. 4.5). The fact that correlations can minimize information content and at the same time maximize robustness highlights the fact that optimizing the amount of information in a population code versus optimizing the ability of that information to be transmitted put very different constraints on neural population codes.

The existence of an optimum where the covariance matrices are matched across layers emphasizes that not all optimally robust population codes are necessarily redundant. (By redundant we mean the population encodes less information than would be encoded by a population of independent cells with the same tuning curves and levels of single neuron trial-to-trial variability [12, 21]; see Fig. 2). Notably, if the effective second layer covariance matrix, $\mathbf{\Sigma}_y$, admits a synergistic population code – wherein more information is encoded in the correlated population versus an uncorrelated one with the same tuning curves and levels of trial-to-trial response variability – then the matched case, $\mathbf{\Sigma}_\xi \propto \mathbf{\Sigma}_y$, will also admit a synergistic population code, and be optimally robust.

Optimally robust, however, does not necessarily mean the majority of the information is transmitted; for that we need another condition. We show in Sec. 4.6 that for non-redundant codes, a large fraction of the information is transmitted only if there are many more neurons in the second layer than in the first. This is typically the case in the periphery. For differential correlations, that condition is not necessary – so long as there are a large number of neurons in both the input and output layers, most of the information is transmitted.

2.5 Nonlinear Gain Functions

So far we have focused on linear gain functions $g(\cdot)$; here we consider nonlinear ones. This case is much harder to analyze, as the effective covariance structure in the second layer, $\mathbf{\Sigma}_{\text{eff},\eta}$, depends on the noise in the first layer (see Methods, Sec. 4.1, especially Eq. (22)). We therefore leave the analysis to Methods (Sec. 4.3); here we briefly summarize the main results. After that we consider two examples of nonlinear gain functions – both involving a thresholding nonlinearity to mimic spike generation.

For linear gain functions we were able to find a whole family of optimal covariance structures, for nonlinear ones we did not even try. Instead, we asked: under what circumstances are differential correlations optimal? Even for this simplified question a definitive answer does not appear to exist. Nevertheless, we can make progress in special cases. When there is no added noise in the second layer (e.g., $\boldsymbol{\eta} = 0$ for the model in Fig. 1B), differential correlations maximize the amount of information that propagates through the nonlinearity, so long as the tuning curves are sufficiently dense relative to the steepness of the tuning curves (meaning that whenever the stimulus changes, the average stimulus-evoked response of at least one neuron also changes; see Methods, Sec. 4.3). If there is added noise at the second layer, differential correlations tend to be optimal in cases where the addition of noise at the first layer, ξ, causes reductions in information, $I_x(s)$. (This means that, so long as there are no *stochastic resonance* effects causing added noise to increase information, then differential correlations are optimal.)

We first check, with simulations, the prediction that differential correlations are optimal if there is no added noise. For that we use a thresholding nonlinearity, chosen for two reasons: it is an extreme nonlinearity, and so should be a strong test of our theory, and it is somewhat realistic in that it mimics spike generation. For this model, the responses at that second layer, y_i, are given by

$$y_i = \Theta(x_i - \theta_i) \tag{11}$$

where Θ is the Heaviside step function ($\Theta(x) = 1$ if $x \geq 0$ and 0 otherwise), and θ_i is the spiking threshold of the i^{th} neuron. This is the popular dichotomized Gaussian model [52–56], which has been shown to provide a good description of population responses in visual cortex, at least in short time windows [54], and to provide high-fidelity descriptions of the responses of integrate-and-fire neurons, again in short time windows [57].

In our simulations with the step function nonlinearity, as for all of the other cases we considered above, the first layer responses are given by the tuning curve plus noise model (Eq. (1)). The tuning curves, $\mathbf{f}(s)$, of the 100-neuron population are again heterogeneous (similar to those in Fig. 2A but with

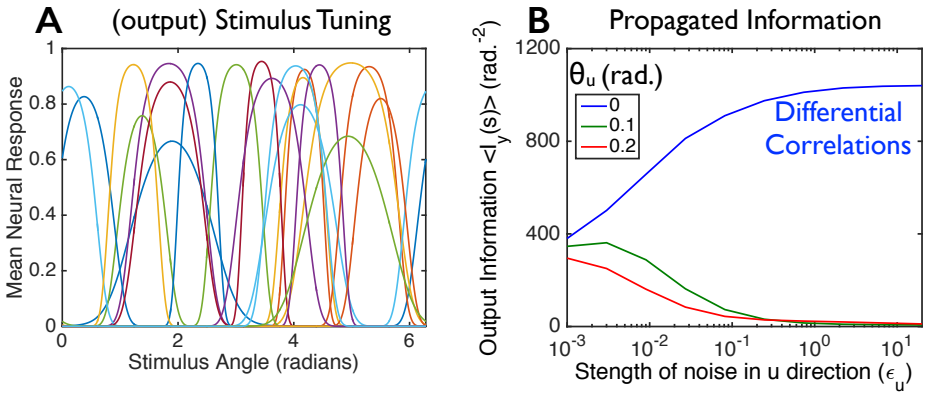

Figure 5. Differential correlations enhance information propagation through "spike-generating" nonlinearities. Responses in the second layer were generated using the dichotomized Gaussian model of spike generation, in which the input from the first layer was simply binarized via a step function (see Eq. (11)). We varied the correlations in these inputs (see Eq. (12)) while keeping the input information and input tuning curves fixed. (**A**) Heterogeneous tuning curves in the second layer, evaluated at $\epsilon_u = 0$; we show a random subset of 20 cells out of the 100-neuron population studied in panel B. (**B**) Information transmitted by the 100-cell spiking population as a function of ϵ_u, which is the strength of the noise in the $\mathbf{u}(s)$ direction, for different angles, θ_u, between $\mathbf{u}$ and $\mathbf{f}'(s)$ (see Eq. (12)). The input information was held fixed as ϵ_u was varied. The information is averaged over 20 evenly spaced stimuli (see Methods, Sec. 4.8.2).

a different random draw from the tuning curve distribution), and the trial-to-trial variability is given by

$$\mathbf{\Sigma}_\xi = \gamma_u \left[\mathbf{\Sigma}_0 + \epsilon_u \mathbf{u}(s)\mathbf{u}(s) \right] \tag{12}$$

with $\mathbf{\Sigma}_0$ given by Eq. (5). This is the same covariance matrix as in Eq. (4b), except that we have included an overall scale factor, γ_u, chosen to ensure that the information in the input layer is independent of both ϵ_u and $\mathbf{u}(s)$ (see Methods, Sec. 4.7, especially Eq. (99)).

Because these (step function) nonlinearities are infinitely steep, the tuning curves are not sufficiently dense for our mathematical analysis to guarantee that differential correlations are optimal for information propagation. However, we argue in Methods, Sec. 4.3, that this should be approximately true for large populations. And indeed, that's what we find with our numerical simulation, as shown in Fig. 5B. When $\theta_u = 0$ (recall that θ_u is the angle between $\mathbf{u}(s)$ and $\mathbf{f}'(s)$), so that $\mathbf{u}(s) = \mathbf{f}'(s)$, the second term in Eq. (12) corresponds to differential correlations; in this case, information increases monotonically with ϵ_u. In other words, information propagated through the step function nonlinearity increases as "upstream" correlations become more like pure differential correlations. In contrast, when θ_u is nonzero (as in Fig. 3A), information does not propagate well: information decreases as ϵ_u increases. This is consistent with our findings for the linear gain function considered in Fig. 2. Thus, differential correlations can optimize information transmission even for a nonlinearity as extreme as a step function.

The lack of explicit added noise at the second layer makes this case somewhat unrealistic. In neural circuits, we expect noise to be added at each stage of processing – if nothing else, due to synaptic failures. We thus considered a model in which noise is added before the spike-generation process,

$$y_i = \Theta(x_i + \zeta_i - \theta_i) \tag{13}$$

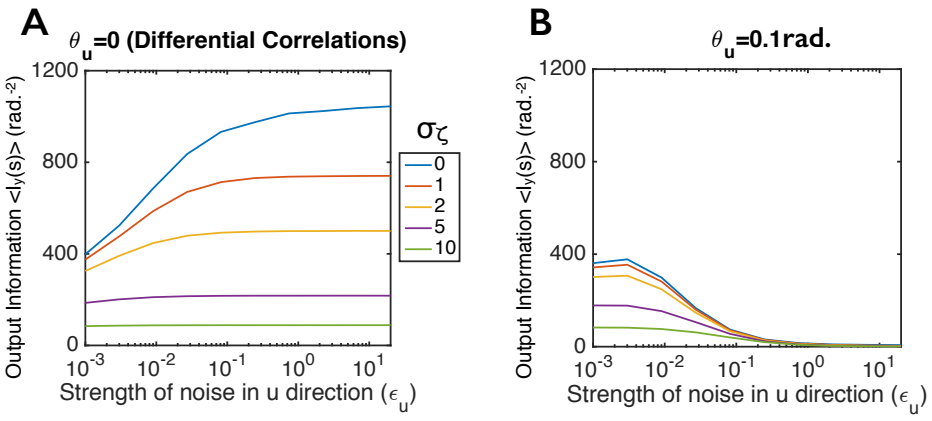

Figure 6. Information propagation through spike-generating nonlinearities with additive input noise. As with Fig. 5, responses in the second layer were generated using the dichotomized model of spike generation, in which the input from the first layer was simply binarized. Here, though, Gaussian noise was added before thresholding; see Eq. (13). We varied the correlations in the input layer (see Eq. (12)) while keeping the input information and input tuning curves fixed for the 100-cell population (same tuning curves and covariance matrices as in Fig. 5). The additive noise at the second layer (the ζ_i) was *iid* Gaussian, with variance σ_ζ^2; different colored lines correspond to different values of σ_ζ^2. (**A**) Output information versus ϵ_u for populations with differential correlations ($\mathbf{u} = \mathbf{f}'(s)$). (**B**) Same as panel A, but for populations that concentrate noise along an axis, $\mathbf{u}$, that makes an angle of 0.1 rad with the $\mathbf{f}'(s)$ direction. For both panels, the input information was held fixed as ϵ_u was varied, and the information was averaged over 20 evenly spaced stimuli.

where ζ_i is zero-mean noise with covariance matrix $\mathbf{\Sigma}_\zeta$.

We computed information for this model using the same input tuning curves, spike thresholds, and covariance matrix, $\mathbf{\Sigma}_\xi$, as without the additional noise (i.e., as in Fig. 5). To mimic the kind of independent noise expected from synaptic failures, we chose the ζ_i to be *iid*, and for simplicity we took them to be Gaussian distributed with variance σ_ζ^2. We computed the amount of stimulus information, $I_y(s)$, for several different levels of the added input noise σ_ζ^2. We found that for all levels of noise, differential correlations increase information transmission ($I_y(s)$ increases monotonically with ϵ_u in Fig. 6A, for which $\theta_u = 0$). And we again found that when the long axis of the covariance ellipse makes a small angle with the signal direction, information propagates poorly (Fig. 6B, for which $\theta_u = 0.1$ rad.).

These numerical findings for a spike-generating nonlinearity with added noise are similar to the previous cases of a linear transfer function, $g(\cdot)$, with added input noise (Figs. 2 and 3), for which we have analytical results, or a spike generating nonlinearity with no added input noise (Fig. 5), for which we do not. We further argue in Methods, Sec. 4.3, that for nonlinear gain functions differential correlations are likely to be optimal if the tuning curves are optimal (in the case of Eq. (13), if the thresholds θ_i are chosen optimally). Taken together, our findings demonstrate that differential correlations in upstream populations generally increase the information that can be propagated downstream through noisy, nonlinear neural circuits.

3 Discussion

Much work in systems neuroscience has investigated the factors that influence the amount of information about a stimulus that is encoded in neural population activity patterns. Here we addressed a related question that is often overlooked: how do correlations between neurons affect the ability of information to propagate robustly through subsequent stages of neural circuitry? The question of robustness is potentially quite important, as the ability of information to propagate determines how much information from the periphery will reach the deeper neural structures that affect decision making and behavior. To investigate this issue, we considered a model with two cell layers. We varied the covariance matrix of the noise in the first layer (while keeping the tuning curves and information in the first layer fixed), and asked how much information could propagate to the second layer. Our main findings were threefold.

First, population codes with different covariance structures but identical tuning curves and equal amounts of encoded information can differ substantially in their robustness to corruption by additional noise (Figs. 2, 5, 6, and 7). Consequently, measurements of information at the sensory periphery are insufficient to understand the ability of those peripheral structures to propagate information to the brain, as that propagation process inevitably adds noise. For instance, populations of independent neurons can be much worse at transmitting information than can populations displaying correlated variability (Fig. 5B). Thus, to understand how the brain efficiently encodes information, we must concern ourselves not just with the amount of information in a population code, but also with the robustness of that encoded information against corruption by noise.

Second, for linear gain functions, or noise-free nonlinear ones with sufficiently dense tuning curves, populations with so-called differential correlations [22] are maximally robust against noise induced by information propagation. This fact may seem surprising given that differential correlations are the only ones that lead to information saturation in large populations [22], and the correlations that minimizme information in general (Methods, Sec. 4.5). However, in hindsight it makes sense: differential correlations correspond to a covariance ellipse aligned with the signal direction (see Fig. 3B), and added noise simply doesn't make it much longer. For nonlinear gain functions combined with arbitrary noise, differential correlations are not guaranteed to yield a globally optimal population code for information propagation. However, for the spike-generating nonlinearity we considered here, differential correlations were at least a local optimum (see Figs. 5 and 6).

Third, while differential correlations optimize robustness, for linear gain functions that optimum is not unique. Instead, there is a continuous family of covariances that exhibit identical robustness to noise (see Fig. 4 and Eq. (7)). However, within this family, only differential correlations yield population codes that are optimally robust independent of the downstream circuitry. Thus, they are the most flexible of the optima: for all other members of the family, the optimal covariance structure in the first layer depends on the noise in subsequent layers, as well as the weights connecting those layers.

The existence of this family of optimal solutions raises an important point with regards to redundancy and robust population coding. Populations with differential correlations – which are among the optimal solutions in terms of robustness – are highly redundant: a population with differential correlations encodes much less information than would be expected from independent populations with the same tuning curves and levels of trial-to-trial variability (Fig. 2C). It is common knowledge that redundancy can enhance robustness of population codes against noise [58], and thus it is worth asking if our robust population coding results are simply an application of this fact. Importantly, the answer is no: as discussed in Sec. 2.4, within the family of optimal correlational structures are codes with minimal redundancy. Moreover, as is shown in Fig. 2B, a code can be redundant without being robust to added noise. In other words, redundancy in a population code is neither necessary, nor sufficient, to ensure that the encoded information is robust against added noise. However, there is an important caveat: unless the number of neurons in the second layer is large relative to the number in the first, and/or the added noise in the second layer is small relative to the noise at the first layer, non-redundant codes tend to lose a large amount of information when corrupted by noise. This contrasts sharply with differential correlations,

which can tolerate large added noise with very little information loss (see Sec. 4.6).

In the case of real neural systems, there will always be a finite amount of information that the population can convey (bounded by the amount of input information that the population receives from upstream sources [59]), and so the question of how best to propagate a (fixed) amount of information is of potentially great relevance for neural communication. Our results suggest that the presence of differential correlations serves to allow population-coded information to propagate robustly. Thus, an observation of these correlations in neural recordings might indicate that the population code is optimized for robustness of the encoded information. At the same time, we note that weak differential correlations might be hard to observe experimentally [22]. Moreover, our calculations indicate that there exists a whole family of possible propagation-enhancing correlation structures, and so differential correlations are not necessary for robust information propagation. This means that observations of either differential correlations, correlation structures matched between subsequent layers of a neural circuit (Fig. 4), or a combination of the above would indicate that the system enables robust information propagation.

How might the nervous system shape its responses so as to generate correlations that enhance information propagation? Recent work identified network mechanisms that can lead to differential correlations [60]. While it is beyond the scope of this work, it would be interesting to explicitly study the network structures that allow encoded information to propagate most robustly through downstream circuits. Relatedly, [38] and [29, 30] asked how the connectivity between layers affects the ability of information to propagate. While we identified the optimal patterns of input to the multi-stage circuit, they identified the optimal anatomy of that circuit itself.

Note that we have used linear Fisher information to quantify the population coding efficacy. Other information measures exist, and it is worth commenting on how much our findings generalize to different measures. In the case of jointly Gaussian stimulus and response distributions, correlations that maximize linear Fisher information also maximize Shannon's mutual information [20]. In that regime our findings should generalize well. Moreover, whenever the neural population response distributions belong to the exponential family with linear sufficient statistics, the linear Fisher information is equivalent to the (nonlinear) "full" Fisher information [29]. In practice, this is a good approximation to primary visual cortical responses to oriented visual stimuli [61, 62], and to other stimulus-evoked responses in other brain areas (see [22] for discussion). Consequently, our use of linear Fisher information in place of other information measures is not a serious limitation.

For encoded sensory information to be useful, it must propagate from the periphery to the deep brain structures that guide behavior. Consequently, information should be encoded in a manner that is robust against corruption that arises during propagation. We showed that the features of population codes that maximize robustness can be substantially different from those that maximize the information content in peripheral layers. Moreover, by elucidating the set of covariances structures that optimize information transmission, we found that redundancy in a population code is neither necessary, nor sufficient, to guarantee robust propagation. In future work, it will be important to determine whether the nervous system uses the class of population codes that maximize information transmission.

Finally, while our main focus was on information propagation, the model we used – linear feedforward weights followed by a nonlinearity – is known to have powerful computational properties [46]. It is, in fact, the basic unit in many deep neural networks. Thus, our main conclusion, which is that differential correlations are typically optimal, applies to any computation that can be performed by this architecture.

4 Methods

Here we provide detailed analysis of the relationship between correlations, feedforward weights, and information propagation. Our methods are organized as follows,

- Sec. 4.1: we derive an expression for the information in the output layer (Eq. (3b)).

- Sec. 4.2: we identify the optimal family of first layer covariance structures when the gain function is linear.
- Sec. 4.3: we consider nonlinear gain functions.
- Sec. 4.4: we take a detailed look at the geometry of information loss.
- Sec. 4.5: we prove that differential correlations minimize information.
- Sec. 4.6: we examine the relationship between coding strategies and information loss.
- Sec. 4.7: we compute information for a noise structure consisting of an arbitrary covariance matrix plus a rank 1 covariance matrix.
- Sec. 4.8: we provide details of our numerical simulations.

4.1 Information in the output layer

Our analysis focuses on information loss through one layer of circuitry; to compute the loss, we need expressions for the linear Fisher information in the first and second layers. Expressions for those two quantities are given in Eqs. (3a) and (3b). The first is standard; here we derive the second.

To make the result as general as possible, we include noise inside the nonlinearity as well as outside it; if nothing else, that's probably a reasonable model for the spiking nonlinearity given in Eq. (13). We thus generalize slightly Eq. (2), and write

$$\mathbf{y} = g(\mathbf{W} \cdot \mathbf{x} + \boldsymbol{\zeta}) + \boldsymbol{\eta} \tag{14}$$

where $\boldsymbol{\zeta}$ is zero mean noise with covariance matrix $\boldsymbol{\Sigma}_\zeta$, and here and in what follows we use the convention that g is a pointwise nonlinearity, so for any vector $\mathbf{v}$, the i^{th} element of $g(\mathbf{v})$ is $g(v_i)$. When $\boldsymbol{\Sigma}_\zeta = 0$, we recover exactly the model in Eq. (2).

Using Eq. (1) for $\mathbf{x}$, Eq. (14) becomes

$$\mathbf{y} = g\big(\mathbf{h}(s) + \mathbf{W} \cdot \boldsymbol{\xi} + \boldsymbol{\zeta}\big) + \boldsymbol{\eta} \tag{15}$$

where, recall, $\boldsymbol{\xi}$ and $\boldsymbol{\eta}$ are zero mean noise with covariance matrices $\boldsymbol{\Sigma}_\xi$ and $\boldsymbol{\Sigma}_\eta$, respectively, and $\mathbf{h}(s)$ is the mean drive to neuron i,

$$\mathbf{h}(s) \equiv \mathbf{W} \cdot \mathbf{f}(s). \tag{16}$$

To compute the linear Fisher information in the second layer, we start with the usual expression,

$$I_y(s) = \frac{\partial \mathrm{E}(\mathbf{y}|s)}{\partial s} \cdot \mathrm{Cov}[\mathbf{y}|s]^{-1} \cdot \frac{\partial \mathrm{E}(\mathbf{y}|s)}{\partial s} \tag{17}$$

where E and Cov denote mean and covariance, respectively. The mean value of $\mathbf{y}$ given s is, via Eq. (15),

$$\mathrm{E}[\mathbf{y}|s] = \mathrm{E}_{\boldsymbol{\xi},\boldsymbol{\zeta}}\Big[g\big(\mathbf{h}(s) + \mathbf{W} \cdot \boldsymbol{\xi} + \boldsymbol{\zeta}\big)\Big] \equiv \overline{g}\big(\mathbf{h}(s)\big). \tag{18}$$

Like $g(\cdot)$, $\overline{g}(\cdot)$ is a pointwise nonlinearity. To compute the covariance, we assume, as in the main text, that $\boldsymbol{\xi}$ and $\boldsymbol{\eta}$ are independent; in addition, we assume that both are independent of $\boldsymbol{\zeta}$. Thus, the covariance

of $\mathbf{y}$ is the sum of the covariances of the first and second terms in Eq. (15). The covariance of the second term is just $\mathbf{\Sigma}_\eta$. The covariance of the first term is harder. To make progress, we start by implicitly defining the quantity $\delta\mathbf{\Sigma}_g(s)$ via

$$\mathrm{Cov}\big[g\big(\mathbf{h}(s)+\mathbf{W}\cdot\boldsymbol{\xi}+\boldsymbol{\zeta}\big)\big] \equiv \delta\mathbf{\Sigma}_g(s) + \Big[\mathbf{W}_{\mathrm{eff}}(s)\cdot\mathbf{\Sigma}_\xi\cdot\mathbf{W}_{\mathrm{eff}}^T(s) + \overline{\mathbf{G}}'(s)\cdot\mathbf{\Sigma}_\zeta(s)\cdot\overline{\mathbf{G}}'(s)\Big] \tag{19}$$

where $\mathbf{W}_{\mathrm{eff}}(s)$ is the actual feedforward weight multiplied by the average slope of g,

$$W_{\mathrm{eff},ij}(s) \equiv \overline{g}'\big(h_i(s)\big)W_{ij}, \tag{20}$$

and $\overline{\mathbf{G}}'(s)$ is the a diagonal matrix with entries corresponding to the average slope of g,

$$\overline{G}'_{ij}(s) \equiv \overline{g}'\big(h_i(s)\big)\delta_{ij}. \tag{21}$$

As in the main text, δ_{ij} is the Kronecker delta and a prime denotes a derivative. The above implicit definition of $\delta\mathbf{\Sigma}_g$ is motivated by the observation that when g is linear, $\delta\mathbf{\Sigma}_g$ vanishes. Below, in Sec. 4.1.1, we show that if $\boldsymbol{\xi}$ is Gaussian, $\delta\mathbf{\Sigma}_g$ is positive semi-definite. Here we assume that the noise is sufficiently close to Gaussian that $\delta\mathbf{\Sigma}_g$ remains positive semi-definite, and thus can be treated as the covariance matrix of an effective noise source. This last assumption is needed in Sec. 4.3, where we argue that information loss is small when $\delta\mathbf{\Sigma}_g$ is small (see text following Eq. (64)).

Making the additional definition

$$\mathbf{\Sigma}_{\mathrm{eff},\eta}(s) \equiv \overline{\mathbf{G}}'(s)\cdot\mathbf{\Sigma}_\zeta(s)\cdot\overline{\mathbf{G}}'(s) + \delta\mathbf{\Sigma}_g(s) + \mathbf{\Sigma}_\eta, \tag{22}$$

and using Eqs. (15) and (19) and the fact that $\boldsymbol{\eta}$ is independent of both $\boldsymbol{\xi}$ and $\boldsymbol{\zeta}$, we see that

$$\mathrm{Cov}[\mathbf{y}|s] = \mathbf{W}_{\mathrm{eff}}(s)\cdot\mathbf{\Sigma}_\xi\cdot\mathbf{W}_{\mathrm{eff}}^T(s) + \mathbf{\Sigma}_{\mathrm{eff},\eta}(s). \tag{23}$$

Combining this with the expression for the mean value of $\mathbf{y}$, Eq. (18), the linear Fisher information, Eq. (17) becomes

$$I_y = \mathbf{f}'\cdot\mathbf{W}_{\mathrm{eff}}^T\cdot\big[\mathbf{W}_{\mathrm{eff}}\cdot\mathbf{\Sigma}_\xi\cdot\mathbf{W}_{\mathrm{eff}}^T + \mathbf{\Sigma}_{\mathrm{eff},\eta}\big]^{-1}\cdot\mathbf{W}_{\mathrm{eff}}\cdot\mathbf{f}' \tag{24}$$

where we used Eqs. (16) and (20) to replace $\partial_s\mathrm{E}(\mathbf{y}|s)$ with $\mathbf{W}_{\mathrm{eff}}\cdot\mathbf{f}'$ and, to reduce clutter, we have suppresed any dependence on s. To pull the effective weights inside the inverse, we use the Woodbury matrix identity to write

$$\begin{aligned}&\mathbf{W}_{\mathrm{eff}}^T\cdot\big[\mathbf{W}_{\mathrm{eff}}\cdot\mathbf{\Sigma}_\xi\cdot\mathbf{W}_{\mathrm{eff}}^T + \mathbf{\Sigma}_{\mathrm{eff},\eta}\big]^{-1}\cdot\mathbf{W}_{\mathrm{eff}} \\ &\quad= \mathbf{W}_{\mathrm{eff}}^T\cdot\mathbf{\Sigma}_{\mathrm{eff},\eta}^{-1}\cdot\mathbf{W}_{\mathrm{eff}} - \mathbf{W}_{\mathrm{eff}}^T\cdot\mathbf{\Sigma}_{\mathrm{eff},\eta}^{-1}\cdot\mathbf{W}_{\mathrm{eff}}\cdot\big[\mathbf{\Sigma}_\xi^{-1} + \mathbf{W}_{\mathrm{eff}}^T\cdot\mathbf{\Sigma}_{\mathrm{eff},\eta}^{-1}\cdot\mathbf{W}_{\mathrm{eff}}\big]^{-1}\cdot\mathbf{W}_{\mathrm{eff}}^T\cdot\mathbf{\Sigma}_{\mathrm{eff},\eta}^{-1}\cdot\mathbf{W}_{\mathrm{eff}}.\end{aligned} \tag{25}$$

Then, using the fact that $[\mathbf{A}+\mathbf{B}]^{-1} = \mathbf{A}^{-1}\cdot[\mathbf{A}^{-1}+\mathbf{B}^{-1}]^{-1}\cdot\mathbf{B}^{-1}$, and applying a very small amount of algebra, this becomes

$$\begin{aligned}&\mathbf{W}_{\mathrm{eff}}^T\cdot\big[\mathbf{W}_{\mathrm{eff}}\cdot\mathbf{\Sigma}_\xi\cdot\mathbf{W}_{\mathrm{eff}}^T + \mathbf{\Sigma}_{\mathrm{eff},\eta}\big]^{-1}\cdot\mathbf{W}_{\mathrm{eff}} \\ &\quad= \mathbf{W}_{\mathrm{eff}}^T\cdot\mathbf{\Sigma}_{\mathrm{eff},\eta}^{-1}\cdot\mathbf{W}_{\mathrm{eff}}\cdot\Big[\mathbf{I} - \mathbf{\Sigma}_\xi\cdot\big[\mathbf{\Sigma}_\xi + (\mathbf{W}_{\mathrm{eff}}^T\cdot\mathbf{\Sigma}_{\mathrm{eff},\eta}^{-1}\cdot\mathbf{W}_{\mathrm{eff}})^{-1}\big]^{-1}\Big]\end{aligned} \tag{26}$$

where $\mathbf{I}$ is the identity matrix. It is then straightforward to show that

$$\mathbf{W}_{\text{eff}}^T \cdot \left[\mathbf{W}_{\text{eff}} \cdot \mathbf{\Sigma}_\xi \cdot \mathbf{W}_{\text{eff}}^T + \mathbf{\Sigma}_{\text{eff},\eta}\right]^{-1} \cdot \mathbf{W}_{\text{eff}} = \left[\mathbf{\Sigma}_\xi + (\mathbf{W}_{\text{eff}}^T \cdot \mathbf{\Sigma}_{\text{eff},\eta}^{-1} \cdot \mathbf{W}_{\text{eff}})^{-1}\right]^{-1}. \tag{27}$$

Inserting this into Eq. (24), we see that the right hand side of that equation is equal to the expression given in Eq. (3b) of the main text.

4.1.1 $\delta\mathbf{\Sigma}_g$ is positive semi-definite for Gaussian noise

To show that $\delta\mathbf{\Sigma}_g$ (defined implicitly in Eq. (19)), is positive semi-definite for Gaussian noise, we'll show that it can be written as a covariance. To simplify the analysis, we make the definition

$$\boldsymbol{\chi} \equiv \mathbf{W} \cdot \boldsymbol{\xi} + \boldsymbol{\zeta}. \tag{28}$$

With this definition,

$$\delta\mathbf{\Sigma}_g = \text{Cov}_{\boldsymbol{\chi}}\left[g\big(\mathbf{h}+\boldsymbol{\chi}\big)\right] - \overline{\mathbf{G}}' \cdot \mathbf{\Sigma}_\chi \cdot \overline{\mathbf{G}}' \tag{29}$$

where here and in what follows we are suppressing the dependence on s, $\mathbf{\Sigma}_\chi$ is the covariance matrix of $\boldsymbol{\chi}$, and $\overline{\mathbf{G}}'$ is defined in Eq. (21). Because we are assuming that both $\boldsymbol{\xi}$ and $\boldsymbol{\zeta}$ are Gaussian, $\boldsymbol{\chi}$ is also Gaussian.

We'll show now that $\delta\mathbf{\Sigma}_g$ is equal to the covariance of the function $g(\mathbf{h}+\boldsymbol{\chi}) - \overline{\mathbf{G}}' \cdot \boldsymbol{\chi}$. We start by noting that

$$\text{Cov}_{\boldsymbol{\chi}}\left[g(\mathbf{h}+\boldsymbol{\chi}) - \overline{\mathbf{G}}' \cdot \boldsymbol{\chi}\right] = \text{Cov}\left[g(\mathbf{h}+\boldsymbol{\chi})\right] - 2\text{Cov}\left[g(\mathbf{h}+\boldsymbol{\chi}), \overline{\mathbf{G}}' \cdot \boldsymbol{\chi}\right] + \overline{\mathbf{G}}' \cdot \mathbf{\Sigma}_\chi \cdot \overline{\mathbf{G}}'. \tag{30}$$

We'll focus on the second term, which is given explicitly by

$$\text{Cov}\left[g(\mathbf{h}+\boldsymbol{\chi}), \overline{\mathbf{G}}' \cdot \boldsymbol{\chi}\right] = \overline{\mathbf{G}}' \cdot \int d\boldsymbol{\chi}\, P(\boldsymbol{\chi})\, \boldsymbol{\chi}\, g(\mathbf{h}+\boldsymbol{\chi}). \tag{31}$$

When $P(\boldsymbol{\chi})$ is Gaussian,

$$P(\boldsymbol{\chi})\boldsymbol{\chi} = -\mathbf{\Sigma}_\chi \cdot \frac{\partial}{\partial \boldsymbol{\chi}} P(\boldsymbol{\chi}). \tag{32}$$

Inserting this into Eq. (31) and integrating by parts, we arrive at

$$\text{Cov}\left[g(\mathbf{h}+\boldsymbol{\chi}), \overline{\mathbf{G}}' \cdot \boldsymbol{\chi}\right] = \overline{\mathbf{G}}' \cdot \mathbf{\Sigma}_\chi \cdot \int d\boldsymbol{\chi}\, P(\boldsymbol{\chi})\, \frac{\partial}{\partial \boldsymbol{\chi}}\, g(\mathbf{h}+\boldsymbol{\chi}). \tag{33}$$

Using the fact that $\partial_{\boldsymbol{\chi}} g(\mathbf{h}+\boldsymbol{\chi}) = \partial_{\mathbf{h}} g(\mathbf{h}+\boldsymbol{\chi})$, the above expression becomes

$$\text{Cov}\left[g(\mathbf{h}+\boldsymbol{\chi}), \overline{\mathbf{G}}' \cdot \boldsymbol{\chi}\right] = \overline{\mathbf{G}}' \cdot \mathbf{\Sigma}_\chi \cdot \frac{\partial}{\partial \mathbf{h}} \int d\boldsymbol{\chi}\, P(\boldsymbol{\chi})\; g(\mathbf{h}+\boldsymbol{\chi}) = \overline{\mathbf{G}}' \cdot \mathbf{\Sigma}_\chi \cdot \overline{\mathbf{G}}'. \tag{34}$$

where the second equality follows from the definition of $\overline{\mathbf{G}}$ (Eq. (21). Inserting this into Eq. (30), we see that the right hand side of Eq. (30) is exactly equal to the right hand side of Eq. (29). Thus, $\delta\mathbf{\Sigma}_g$ can be written as a covariance, and so it must be positive semi-definite.

4.2 Identifying the family of optimal covariance matrices

Here we address the question: what noise covariance matrix optimizes information transmission? In other words, what covariance matrix $\mathbf{\Sigma}_\xi$ maximizes the information given in Eq. (3b)? That is hard to answer when g is nonlinear, because in that case $\mathbf{\Sigma}_{\text{eff},\eta}$ depends on $\mathbf{\Sigma}_\xi$ via $\delta\mathbf{\Sigma}_g$ (see Eqs. (19) and (22)). In this section, then, we consider linear gain functions; in the next we consider nonlinear ones. To make our expressions more readable, we generally suppress the dependence on s.

Our goal is to maximize I_y with I_x fixed. Using the definition of $\mathbf{\Sigma}_y$ given in Eq. (8), for linear gain functions the information in the second layer (Eq. (3b)) is written

$$I_y = \mathbf{f}' \cdot \left[\mathbf{\Sigma}_\xi + \mathbf{\Sigma}_y\right]^{-1} \cdot \mathbf{f}'. \tag{35}$$

We use Lagrange multipliers,

$$\frac{\partial}{\partial \mathbf{\Sigma}_\xi}\Big[\mathbf{f}' \cdot [\mathbf{\Sigma}_\xi + \mathbf{\Sigma}_y]^{-1} \cdot \mathbf{f}' - \lambda\big(\mathbf{f}' \cdot \mathbf{\Sigma}_\xi^{-1} \cdot \mathbf{f}' - I_x\big)\Big] = 0, \tag{36}$$

where λ is a Lagrange multiplier that enforces the constraint $\mathbf{f}' \cdot \mathbf{\Sigma}_\xi^{-1} \cdot \mathbf{f}' = I_x$. Taking the derivative and setting it to zero yields

$$[\mathbf{\Sigma}_\xi + \mathbf{\Sigma}_y]^{-1} \cdot \mathbf{f}'\mathbf{f}' \cdot [\mathbf{\Sigma}_\xi + \mathbf{\Sigma}_y]^{-1} = \lambda \mathbf{\Sigma}_\xi^{-1} \cdot \mathbf{f}'\mathbf{f}' \cdot \mathbf{\Sigma}_\xi^{-1}\,. \tag{37}$$

In deriving this expression we used the fact that the gain functions are linear, which implies that $\mathbf{\Sigma}_y$ does not depend on $\mathbf{\Sigma}_\xi$. Multiplying by $\mathbf{\Sigma}_\xi + \mathbf{\Sigma}_y$ on both the left and right, we arrive at

$$\mathbf{f}'\mathbf{f}' = \lambda[\mathbf{I} + \mathbf{\Sigma}_y \cdot \mathbf{\Sigma}_\xi^{-1}] \cdot \mathbf{f}'\mathbf{f}' \cdot [\mathbf{I} + \mathbf{\Sigma}_\xi^{-1} \cdot \mathbf{\Sigma}_y]. \tag{38}$$

This is satisfied when

$$\mathbf{\Sigma}_y \cdot \mathbf{\Sigma}_\xi^{-1} \cdot \mathbf{f}' \propto \mathbf{f}'\,. \tag{39}$$

There are two ways this can happen,

$$\mathbf{\Sigma}_y \cdot \mathbf{\Sigma}_\xi^{-1} \propto \mathbf{I} \tag{40a}$$

$$\mathbf{\Sigma}_y \cdot \mathbf{\Sigma}_\xi^{-1} \propto \mathbf{f}'\mathbf{a} \tag{40b}$$

where $\mathbf{a}$ is an arbitrary vector. Combining these linearly, taking into account that $\mathbf{\Sigma}_\xi$ is a covariance matrix and thus symmetric, and enforcing equality in Eq. (38), we arrive at

$$\mathbf{\Sigma}_\xi^{-1} = \frac{I_x}{\alpha I_\eta}\mathbf{\Sigma}_y^{-1} + \frac{(\alpha-1)I_x}{\alpha I_\eta^2}\mathbf{\Sigma}_y^{-1} \cdot \mathbf{f}'\mathbf{f}' \cdot \mathbf{\Sigma}_y^{-1} + \mathbf{P} \cdot \mathbf{\Omega}^{-1} \cdot \mathbf{P} \tag{41}$$

where I_η is the information the output layer would have if there was no noise in the input layer,

$$I_\eta(s) = \mathbf{f}'(s) \cdot \mathbf{\Sigma}_y^{-1} \cdot \mathbf{f}'(s) \tag{42}$$

(this is the same expression as in Eq. (9), it's repeated here for convenience), $\mathbf{\Omega}$ is an arbitrary symmetric matrix, $\mathbf{P}$ is a projection operator, chosen so that $\mathbf{P} \cdot \mathbf{f}' = 0$,

$$\mathbf{P} \equiv \mathbf{I} - \frac{\mathbf{f}'\mathbf{f}'}{\mathbf{f}' \cdot \mathbf{f}'}\,, \tag{43}$$

and α is arbitrary (but subject to the constraint that $\mathbf{\Sigma}_\xi$ has no negative eigenvalues). Note that $\mathbf{P}$ is a linear combination of the right hand sides of Eqs. (40a) and (40b)), with $\mathbf{a} = \mathbf{f}'$ in the latter equation. It is straightforward to verify that when $\mathbf{\Sigma}_\xi$ is given by Eq. (41), Eq. (38) is satisfied.

To find an explicit expression for $\mathbf{\Sigma}_\xi$, not just its inverse, we apply the Woodbury matrix identity to Eq. (41); that gives us

$$\mathbf{\Sigma}_\xi = \mathbf{\Sigma}_\alpha - \mathbf{\Sigma}_\alpha \cdot \mathbf{P} \cdot \left(\mathbf{\Omega} + \mathbf{P} \cdot \mathbf{\Sigma}_\alpha \cdot \mathbf{P}\right)^{-1} \cdot \mathbf{P} \cdot \mathbf{\Sigma}_\alpha \tag{44}$$

where

$$\mathbf{\Sigma}_\alpha \equiv \left[\frac{I_x}{\alpha I_\eta}\mathbf{\Sigma}_y^{-1} + \frac{(\alpha-1)I_x}{\alpha I_\eta^2}\mathbf{\Sigma}_y^{-1} \cdot \mathbf{f}'\mathbf{f}' \cdot \mathbf{\Sigma}_y^{-1}\right]^{-1} = \frac{\alpha I_\eta}{I_x}\left[\mathbf{\Sigma}_y - \frac{(\alpha-1)\mathbf{f}'\mathbf{f}'}{\alpha I_\eta}\right]. \tag{45}$$

Inserting this into Eq. (44), we arrive at

$$\mathbf{\Sigma}_\xi = \frac{\alpha I_\eta \mathbf{\Sigma}_y}{I_x} + \frac{(1-\alpha)\mathbf{f}'\mathbf{f}'}{I_x} - \left(\frac{\alpha I_\eta}{I_x}\right)^2 \mathbf{\Sigma}_y \cdot \mathbf{P} \cdot \left(\mathbf{\Omega} + \frac{\alpha I_\eta}{I_x}\mathbf{P} \cdot \mathbf{\Sigma}_y \cdot \mathbf{P}\right)^{-1} \cdot \mathbf{P} \cdot \mathbf{\Sigma}_y \,. \tag{46}$$

This is the same as Eq. (7) in the main text, except in that equation we let $\mathbf{\Omega}$ go to ∞, so we ignore the projection-related term. Ignoring that term is reasonable, as it just puts noise in a direction perpendicular to $\mathbf{f}'$, and so has no effect on the information.

By choosing different scalars α and matrices $\mathbf{\Omega}$, a family of optimal $\mathbf{\Sigma}_\xi$ is obtained. These all have the same input information, I_x, and the same output information, I_y, after corruption by noise. An especially interesting covariance matrix is found in the limit $\alpha = 0$, in which case

$$\mathbf{\Sigma}_\xi = \frac{\mathbf{f}'\mathbf{f}'}{I_x} \,. \tag{47}$$

These are so-called differential correlations [22]. Importantly, the choice $\alpha = 0$ is the only one for which the optimal correlational structure is independent of the correlations in the output layer, $\mathbf{\Sigma}_y$. Note that pure differential correlations don't satisfy Eq. (39). As such, they represent a singular limit, in the sense that $\mathbf{\Sigma}_\xi$ in Eq. (46) satisfies Eq. (39) with alpha arbitrarily small, but not precisely zero.

The other covariance that we highlight in the text is found for $\alpha = 1$ and $\mathbf{\Omega} \to \infty$, in which case $\mathbf{\Sigma}_\xi = \frac{\alpha I_\eta}{I_x}\mathbf{\Sigma}_y$. This is the matched covariance case.

4.3 Nonlinear gain functions

We now focus on differential correlations, and determine conditions under which they are optimal for information propagation when the gain function, $g(\cdot)$, is nonlinear. In this regime, the effective noise in the second layer (the second term in brackets in Eq. (3b)) depends on $\mathbf{\Sigma}_\xi$. This greatly complicates the analysis, and to make headway we need to reformulate our mathematical description of differential correlations. This reformulation is based on the observation that differential correlations correspond to trial-to-trial variability in the value of the stimulus, s [22]. Consequently, the encoding model in the input layer can be written as a multi-step process,

$$s = s_0 + \delta_s \tag{48a}$$
$$\mathbf{x} = \mathbf{f}(s) + \boldsymbol{\xi}(s) \tag{48b}$$
$$\mathbf{y} = g\big(\mathbf{W} \cdot \mathbf{x}(s) + \boldsymbol{\zeta}\big) + \boldsymbol{\eta}(s). \tag{48c}$$

Here s_0 is the value of the stimulus that is actually presented. However, the neurons in the input layer, $\mathbf{x}$, encode s – a corrupted version of s_0. This is indicated by Eq. (48a), which tells us that s deviates on a trial-to-trial basis from s_0, with deviations that are described by a zero-mean random variable, δ_s.

To see that this model does indeed exhibit differential correlations, we Taylor expand Eq. (48b) around s_0, yielding a model of the form

$$\mathbf{x} \approx \mathbf{f}(s_0) + \mathbf{f}'(s_0)\delta_s + \boldsymbol{\xi}(s_0), \tag{49}$$

for which the covariance matrix is

$$\mathrm{Cov}[\mathbf{x}] = \mathrm{Var}[\delta_s]\mathbf{f}'(s_0)\mathbf{f}'(s_0) + \mathrm{Cov}[\boldsymbol{\xi}|s_0]. \tag{50}$$

The first term corresponds to differential correlations.

Equations (48b) and Eq. (48c) correspond exactly to our previous model (Eq. (4)a). Consequently, the information about s in the first and second layers are still given by Eqs. (3a) and (3b) of the main text. However, we can't use those equations for the information about s_0. For that, we focus on the variance of its optimal estimator given $\mathbf{x}$, which we denote $\hat{s}_0$. Because of the Markov structure of our model ($s_0 \leftrightarrow s \leftrightarrow x$), we can construct $\hat{s}_0$ by first considering the optimal estimator of s_0 given s, and then the optimal estimator of s given $\mathbf{x}$. The variance of $\hat{s}_0$ given $\mathbf{x}$ is then simply the sum of the variances of these two (independent) noise sources.

The optimal estimator of s_0 given s is simply s, with conditional variance $\mathrm{Var}[\hat{s}_0(s)|s_0] = \mathrm{Var}[\delta_s]$. The optimal estimator of s given $\mathbf{x}$ is $\hat{s}(\mathbf{x})$, with variance $\mathrm{Var}[\hat{s}(\mathbf{x})|s]$. Consequently,

$$\mathrm{Var}[\hat{s}_0|s_0] = \mathrm{Var}[\delta_s] + \int ds\, P(s|s_0)\mathrm{Var}[\hat{s}(\mathbf{x})|s]. \tag{51}$$

As usual, we approximate the variance of $\hat{s}(\mathbf{x})$ given s by the linear Fisher information, yielding an approximation for the total Fisher information about s_0 given $\mathbf{x}$,

$$\frac{1}{I_x^{\mathrm{tot}}(s_0)} = \mathrm{Var}[\delta_s] + \int ds\, \frac{P(s|s_0)}{I_x(s)}\,. \tag{52}$$

Similarly, the Fisher information about s_0 given $\mathbf{y}$ is approximated by

$$\frac{1}{I_y^{\mathrm{tot}}(s_0)} = \mathrm{Var}[\delta_s] + \int ds\, \frac{P(s|s_0)}{I_y(s)}\,. \tag{53}$$

Note that we are slightly abusing notation here: above, $I_x(s)$ and $I_y(s)$ referred to the total information about the stimulus; now they refer to the information about the stimulus that is encoded in the first layer, which is different from the actual stimulus, s_0. However, it is a convenient abuse, as it allows us to take over our previous results without introducing much new notation.

Our first step is to parametrize the covariance matrix, $\boldsymbol{\xi}$, and $\mathrm{Var}[\delta_s]$, in a way that ensures that the information in the first layer $I_x^{\mathrm{tot}}(s_0)$ remains fixed while we vary $\boldsymbol{\xi}$ and $\mathrm{Var}[\delta_s]$. A convenient choice is

$$\mathrm{Var}[\delta_s] = \frac{1}{I_x^{\mathrm{tot}}} \int ds\, P(s|s_0) \frac{\epsilon I_0(s)}{1+\epsilon I_0(s)} \tag{54a}$$

$$\boldsymbol{\Sigma}_\xi(s) = \frac{1}{I_x^{\mathrm{tot}}} \frac{I_0(s)\boldsymbol{\Sigma}_0(s)}{1+\epsilon I_0(s)}\,, \tag{54b}$$

where

$$I_0(s) \equiv \mathbf{f}'(s) \cdot \mathbf{\Sigma}_0^{-1}(s) \cdot \mathbf{f}'(s). \tag{55}$$

Inserting Eq. (54) into Eq. (52), we see that $I_x^{\rm tot}(s_0) = I_x^{\rm tot}$, independent of $\mathbf{\Sigma}_0(s)$.

The information in the second layer about s, $I_y(s)$, is given by Eq. (3b), with $\mathbf{\Sigma}_{{\rm eff},\eta}$ given in Eq. (22). It is convenient to make the definition

$$\mathbf{\Sigma}_{{\rm eff},y} \equiv (\mathbf{W}_{\rm eff}^T \cdot \mathbf{\Sigma}_{{\rm eff},\eta}^{-1} \cdot \mathbf{W}_{\rm eff})^{-1}. \tag{56}$$

This is the analog of Eq. (8), but for nonlinear gain functions. It is clear from Eqs. (22) and (19) that $\mathbf{\Sigma}_{{\rm eff},\eta}$ depends on $\mathbf{\Sigma}_\xi$; consequently, it depends on ϵ.

To maximize information with respect to ϵ, we take a two step approach. We write

$$I_y(s;\epsilon,\epsilon_0) \equiv \mathbf{f}'^T(s)\left[\mathbf{\Sigma}_\xi(s,\epsilon) + \mathbf{\Sigma}_{{\rm eff},y}(s,\epsilon_0)\right]^{-1}\mathbf{f}'(s). \tag{57}$$

Here $\mathbf{\Sigma}_\xi(s,\epsilon)$ and $\mathbf{\Sigma}_y(s,\epsilon_0)$ are the same as in Eqs. (54b) and (56); we have just made the dependence on ϵ explicit. The two steps are to maximize first with respect to ϵ, then with respect to ϵ_0. If the two maxima occurr in the same place, then we have identified the covariance structure that optimizes information transmission.

In the first step we differentiate $I_y^{\rm tot}(s;\epsilon,\epsilon_0)$ with respect to ϵ. To simplify the expressions, we make the definition

$$\mathbf{\Sigma}_{\rm tot}(s,\epsilon,\epsilon_0) \equiv \mathbf{\Sigma}_\xi(s,\epsilon) + \mathbf{\Sigma}_{{\rm eff},y}(s,\epsilon_0). \tag{58}$$

Combining Eqs. (53), (54) and (57), we have

$$\frac{\partial}{\partial\epsilon}\frac{1}{I_y^{\rm tot}(s_0;\epsilon,\epsilon_0)} = \frac{1}{I_x^{\rm tot}}\int ds\, P(s|s_0)\frac{I_0}{(1+\epsilon I_0)^2} + \int ds\,\frac{P(s|s_0)}{I_y^2}\,\mathbf{f}'\cdot\mathbf{\Sigma}_{\rm tot}^{-1}\cdot\frac{\partial\mathbf{\Sigma}_\xi(s,\epsilon)}{\partial\epsilon}\cdot\mathbf{\Sigma}_{\rm tot}^{-1}\cdot\mathbf{f}' \tag{59}$$

where we used the fact that for any square matrix $\mathbf{A}(x)$, $(d/dx)\mathbf{A}^{-1} = -\mathbf{A}^{-1}\cdot d\mathbf{A}/dx\cdot\mathbf{A}^{-1}$, and we suppressed much of the s, ϵ and ϵ_0 dependence for clarity. Using Eq. (54b) for $\mathbf{\Sigma}_\xi(s,\epsilon)$, the derivative with respect to ϵ in the second term is straightforward,

$$\begin{aligned}\frac{\partial}{\partial\epsilon}\frac{1}{I_y^{\rm tot}(s_0;\epsilon,\epsilon_0)} &= \frac{1}{I_x^{\rm tot}}\int ds\, P(s|s_0)\frac{I_0}{(1+\epsilon I_0)^2} - \int ds\,\frac{P(s|s_0)}{I_y^2}\,\mathbf{f}'\cdot\mathbf{\Sigma}_{\rm tot}^{-1}\cdot\frac{I_0^2\mathbf{\Sigma}_0}{I_x^{\rm tot}(1+\epsilon I_0)^2}\cdot\mathbf{\Sigma}_{\rm tot}^{-1}\cdot\mathbf{f}' \\ &= \frac{1}{I_x^{\rm tot}}\int ds\, P(s|s_0)\frac{I_0^2}{I_y^2(1+\epsilon I_0)^2}\left[\frac{I_y^2}{I_0} - \mathbf{f}'\cdot\mathbf{\Sigma}_{\rm tot}^{-1}\cdot\mathbf{\Sigma}_0\cdot\mathbf{\Sigma}_{\rm tot}^{-1}\cdot\mathbf{f}'\right]\end{aligned} \tag{60}$$

Then, applying the definition $I_y(s) = \mathbf{f}'\cdot\mathbf{\Sigma}_{\rm tot}^{-1}\cdot\mathbf{f}'$ (see Eqs. (57) and (58)), and making the new definition

$$\mathbf{V} \equiv \mathbf{f}'\cdot\mathbf{\Sigma}_{\rm tot}^{-1}\cdot\mathbf{\Sigma}_0^{1/2}, \tag{61}$$

we arrive at the expression

$$\frac{\partial}{\partial\epsilon}\frac{1}{I_y^{\rm tot}(s_0;\epsilon,\epsilon_0)} = \frac{1}{I_x^{\rm tot}}\int ds\, P(s|s_0)\,\frac{I_0^2}{I_y^2(1+\epsilon I_0)^2}\,\mathbf{V}\cdot\left[\frac{\boldsymbol{\Sigma}_0^{-1/2}\mathbf{f}'\mathbf{f}'^T\boldsymbol{\Sigma}_0^{-1/2}}{I_0} - \mathbf{I}\right]\cdot\mathbf{V}. \tag{62}$$

The right hand side of Eq. (62) is negative or zero if the term in brackets is negative semi-definite; that is, if all its eigenvalues are non-positive. Since the term in square brackets is a rank one matrix minus the identity, all but one of its eigenvalues are equal to -1. The remaining eigenvalue is 0, with corresponding eigenvector $\boldsymbol{\Sigma}_0^{-1/2}\cdot\mathbf{f}'$ (see Eq. (55)). Thus, $\partial(1/I_y^{\rm tot}(s_0;\epsilon,\epsilon_0)/\partial\epsilon) \le 0$, and $I_y^{\rm tot}(s_0;\epsilon,\epsilon_0)$ must have a global maximum at $\epsilon=\infty$. If g is linear, $\boldsymbol{\Sigma}_{{\rm eff},y}$ doesn't depend on ϵ_0, and $\epsilon=\infty$ corresponds to pure differential correlations. We have, therefore, recovered the $\alpha=0$ limit of Eq. (46).

When $\epsilon=\infty$, $\boldsymbol{\Sigma}_\xi$ vanishes, and so the expression for the information in the second layer simplifies considerably. Combining Eqs. (53) and (54a), we have, in the $\epsilon\to\infty$ limit,

$$\frac{1}{I_y^{\rm tot}(s_0;\infty,\epsilon_0)} = \frac{1}{I_x^{\rm tot}} + \int ds\,\frac{P(s|s_0)}{I_y(s;\infty,\epsilon_0)} \tag{63}$$

where

$$I_y(s;\infty,\epsilon_0) = \mathbf{f}'(s)\cdot\mathbf{W}_{\rm eff}^T(s;\epsilon_0)\cdot\left[\boldsymbol{\Sigma}_\eta + \overline{\mathbf{G}}'(s)\cdot\boldsymbol{\Sigma}_\zeta(s)\cdot\overline{\mathbf{G}}'(s) + \delta\boldsymbol{\Sigma}_g(s;\epsilon_0)\right]^{-1}\cdot\mathbf{W}_{\rm eff}(s;\epsilon_0)\cdot\mathbf{f}'(s). \tag{64}$$

The latter equation follows by combining the fact that $\boldsymbol{\Sigma}_\xi(s,\infty)=0$ (Eq. (54b)) with the definitions of $\boldsymbol{\Sigma}_{{\rm eff},y}$ and $\boldsymbol{\Sigma}_{{\rm eff},\eta}$ (Eqs. (56) and (22), respectively).

The total information in the output layer is maximized when $I_y(s_0;\infty,\epsilon_0)$ is maximized. That quantity depends on ϵ_0 via $\boldsymbol{\Sigma}_\xi(s,\epsilon_0)$, the noise covariance in the input layer. As can be seen from Eq. (54b), larger ϵ_0 implies smaller $\boldsymbol{\Sigma}_\xi(s,\epsilon_0)$. That has two effects. First, when $\boldsymbol{\Sigma}_\xi(s,\epsilon_0)$ is small enough, the covariance matrix $\delta\boldsymbol{\Sigma}_g$ becomes small (see Eq. (19), and note that $\delta\boldsymbol{\Sigma}_g$ is positive definite, as shown in Sec. 4.1.1). This tends to increase $I_y(s)$. However, the effective tuning curves, $\mathbf{W}_{\rm eff}(s;\epsilon)\cdot\mathbf{f}(s)$, also depend on $\boldsymbol{\Sigma}_\xi(s,\epsilon_0)$ (see Eq. (20)). It is possible that increasing $\boldsymbol{\Sigma}_\xi(s,\epsilon_0)$ modifes the tuning curves such that $I_y(s)$ increases. Consequently, it is impossible to make completely general statements.

Nevertheless, we can identify two regimes. First, if there is no added noise in the output layer ($\boldsymbol{\eta}=\boldsymbol{\zeta}=0$), then $I_y(s;\infty,\epsilon)$ goes to ∞ as ϵ_0 goes to ∞, thus maximizing the total information. This holds, however, only if the tuning curves are sufficiently dense relative to the steepness of the tuning curves; otherwise, the Fisher information is no longer a good approximation to the true information. For smooth tuning curves this is generally satisfied, but it is not satisfied for the noise-free spike generating mechanism we consider in the main text (Eq. (11)), since for that nonlinearity $\mathbf{f}'(s)=0$ with probability 1. We expect, though, that in the absence of noise, this particular nonlinearity introduces an error that is $\mathcal{O}(1/n)$, implying that $I_y(s;\infty,\epsilon)\propto n^2$. Numerical simulations (not shown) corroborated this scaling. Thus, for sufficiently large populations, differential correlations are optimal for the noise-free spike-generating nonlinearity. Note, though, that the thresholds must be chosen so that there are always both active and silent neurons; otherwise, in the limit that $\boldsymbol{\Sigma}_\xi$ vanishes, the activity will contain no information at all about the stimulus.

The second regime is one in which the tuning curves have been optimized. In this case, modifying the tuning curves by adding noise decreases information, and again differential correlations optimize information transmission.

To summarize, we have analyzed the scenario considered in Sec. 2.5 – namely, the neural activities at the second layer, $\mathbf{y}$, are given by a nonlinear function of the neural activities at the first layer, $\mathbf{x}$, with noise added both before and after the nonlinearity. In this case, whether or not differential correlations in the first layer optimize information transmission depends on the details. They do if g is linear, the tuning

curves are optimal, or there is no added noise in the second layer and the tuning curves are sufficiently dense relative to the steepness of the tuning curves. If none of these are satisfied, however, differential correlations may be sub-optimal.

4.4 Analysis behind the geometry of information loss

Our goal in this section is to make more rigorous the geometrical arguments in Fig. 3. We start with the observation that, for Gaussian distributed neural responses, the 1 standard-deviation probability contours for the responses in the first layer (magenta ellipses in Fig. 3) are defined by

$$\Delta\mathbf{r} \cdot \mathbf{\Sigma}_\xi^{-1} \cdot \Delta\mathbf{r} = 1, \tag{65}$$

where $\Delta\mathbf{r} \equiv \mathbf{f}(s) - \mathbf{r}$ represents fluctuations around the mean response to stimulus s. In two dimensions, which we'll focus on here, Eq. (65) becomes

$$\frac{\Delta r_1^2}{\sigma_1^2} + \frac{\Delta r_2^2}{\sigma_2^2} = 1 \tag{66}$$

where σ_1 and σ_2 are the lengths of the principal axes of the covariance ellipse (so σ_1^2 and σ_2^2 are the eigenvalues of $\mathbf{\Sigma}_\xi$) and Δr_1 and Δr_2 are distances spanned by the magenta ellipses along those axes.

As shown in Fig. 3, the intersection between the magenta ellipse (the one defined in Eq. (66)) and the signal curve tells us the uncertainty in the value of the stimulus. To quantify this uncertainty, we simply set $\Delta\mathbf{r}$ to $\mathbf{f}'(s)\Delta s_x$ (the subscript x indicates that this is the uncertainty in the input layer), insert that into Eq. (66), and solve for Δs_x. Defining θ to be the angle between $\mathbf{f}'(s)$ and the long principal axis (see Fig. 3, and note that $\theta = 0$ in panel B), and letting σ_1 correspond to the length of the ellipse's major axis (so $\sigma_1 > \sigma_2$), we have

$$|\mathbf{f}'(s)|^2 \left[\frac{\cos^2\theta}{\sigma_1^2} + \frac{\sin^2\theta}{\sigma_2^2}\right] = \frac{1}{\Delta s_x^2}. \tag{67}$$

The left hand side is the linear Fisher information in the first layer [10], a fact that is useful primarily because it validates our (relatively informal) derivation. More importantly, we can now see how *iid* noise affects information. The addition of *iid* noise simply increases the eigenvalues by σ^2, so the ratio of the information in the output layer to that in the input layer is

$$\frac{I_y}{I_x} = \frac{\Delta s_x^2}{\Delta s_y^2} = \frac{\frac{\cos^2\theta}{\sigma_1^2+\sigma^2} + \frac{\sin^2\theta}{\sigma_2^2+\sigma^2}}{\frac{\cos^2\theta}{\sigma_1^2} + \frac{\sin^2\theta}{\sigma_2^2}}. \tag{68}$$

We can identify two limits. First, if $\theta = 0$ (as it is in Fig. 3B), this ratio reduces to

$$\left.\frac{I_y}{I_x}\right|_{\theta=0} = \frac{\sigma_1^2}{\sigma_1^2 + \sigma^2}\,. \tag{69}$$

Second, if $\tan\theta \gg \sigma_2/\sigma_1$ (which essentially means the green line in Fig. 3 intersects the covariance ellipse on the side, as in panel A, rather than somewhere near the end, as in panel B), the ratio of the informations becomes

$$\left.\frac{I_y}{I_x}\right|_{\tan\theta \gg \sigma_2/\sigma_1} \approx \frac{\sigma_2^2}{\sigma_2^2+\sigma^2}\,. \tag{70}$$

Because $\sigma_1 > \sigma_2$, the information loss is larger in the second case than in the first. And the longer and skinnier the covariance ellipse, the larger the difference in information loss. Thus, this analysis quantifies the geometrical picture given in Fig. 3, in which there is larger information loss in panel A (where $\theta > 0$) than in panel B (where $\theta = 0$).

4.5 Minimum information

Here we ask: what correlational structure minimizes linear Fisher information? To answer that, we use the multi-dimensional analog of Eq. (67),

$$I_x(s) = |\mathbf{f}'(s)|^2 \sum_k \frac{\cos^2\theta_k}{\sigma_k^2} \tag{71}$$

where σ_k^2 is the k^{th} eigenvalue of the noise covariance matrix and θ_k is the angle between $\mathbf{f}'(s)$ and the k^{th} eigenvector [10]. We would like to minimize $I_x(s)$ with respect to the angles, θ_k, and the eigenvalues, σ_k^2. Without constraints, this problem is trivial: information is minimized by having infinite variances for the neural activities. To make the problem better-formulated, we add a constraint that prevents the optimization procedure from simply identifying that trivial solution.

We'll come to the constraint shortly, but first we'll minimize information with respect to the angles, θ_k. That minimum occurs when the eigenvector corresponding to the largest eigenvalue is parallel to $\mathbf{f}'(s)$; ordering the eigenvalues so that σ_0^2 is the largest eigenvalue, we have $\cos\theta_0 = 1$ and $\cos\theta_{k>0} = 0$. Consequently, the information at the minimum is

$$I_x(s) = \frac{|\mathbf{f}'(s)|^2}{\sigma_0^2}. \tag{72}$$

The next step is to minimize $I_x(s)$ with respect to the eigenvalues, subject to a constraint on the covariance matrix. We consider constraints of the form

$$C(\sigma_0^2, \sigma_1^2, ...) \le C_0 \tag{73}$$

where, to avoid the trivial solution (of infinite neural variances), C is an increasing function of each of it's arguments: for all k,

$$\frac{\partial C(\sigma_0^2, \sigma_1^2, ...)}{\partial \sigma_k^2} \ge 0. \tag{74}$$

Examples of $C(\sigma_0^2, \sigma_1^2, ...)$ are the trace of the covariance matrix (the sum of the eigenvalues) and the Frobenius norm (the square root of the sum of the squares of the eigenvalues).

Because of Eq. (74), the information, Eq. (72), is minimized and the constraint, Eq. (73), is satisfied when all the eigenvalues except σ_0^2 are zero. At this global minimum, the covariance matrix, $\mathbf{\Sigma}_\xi$, displays purely differential correlations,

$$\mathbf{\Sigma}_\xi = \sigma_0^2 \mathbf{v}_0 \mathbf{v}_0 \propto \mathbf{f}'(s)\mathbf{f}'(s) \tag{75}$$

where $\mathbf{v}_0$ is the eigenvector associated with the largest eigenvalue. The last term in this expression follows because the above minimization with respect to the angles forced $\mathbf{v}_0$ to be parallel to $\mathbf{f}'(s)$. Thus, for a broad, and reasonable, class of constraints on the covariance matrix, differential correlations minimize information.

4.6 Variances of neural responses, and robustness to added noise, for different coding strategies

Throughout most of our analysis we focused on optimality of information transmission. However, also important is how much information is transmitted at the optimum. That's the subject of this section. For simplicity we consider a linear gain function, which we set, without loss of generality, to the identity. That allows us to use the analysis in Sec. 4.2, and in particular Eq. (46), which links the noise in the input and output layers.

Our starting point is the derivation of an expression for the ratio of the information in the output layer to that in the input layer. To do that, we dot both sides of Eq. (37) by $\mathbf{f}'$ on the left and right sides and solve for λ; we then do the same, except we dot with $\mathbf{f}' \cdot \mathbf{\Sigma}_y^{-1} \cdot [\mathbf{\Sigma}_\xi + \mathbf{\Sigma}_y]$ on the left and its transpose on the right. This yields, after a small amount of algebra,

$$\frac{I_y}{I_x} = \frac{I_\eta}{I_\eta + I_x} = \frac{1}{1 + I_x/I_\eta} \tag{76}$$

where I_x, I_y and I_η are given by Eqs. (3a), (3b) and (9), respectively. For information to be transmitted efficiently, I_x, the information in the input layer, must be small compared to I_η, the information associated with the added noise in the output layer. Below, we investigate the conditions under which $I_x \ll I_\eta$, and thus when information loss is small.

Our strategy is to express I_x/I_η in terms of the single neuron variability, quantified as the average variance – something that has an easy interpretation. We consider two cases: the weights are set to the identity ($\mathbf{W} = \mathbf{I}$), and the weights are more realistic (each neuron in the input layer connects to a large number of neurons in the output layer). The first case, identity weights, is not very realistic; we include it because it is much simpler than the second.

While the analysis is straightforward, it is somewhat heavy on the algebra, so we summarize the results here. We consider two extremes in the family of optimal covariance structures: the "matched" case ($\alpha = 1$ in Eq. (46), and, for simplicity, $\Omega = \infty$) and differential correlations ($\alpha = 0$). For matched covariances, near complete information transfer ($I_x \ll I_\eta$) requires the effective variance of the noise in the second layer to be small. For identity feedforward weights, the effective variance in the input and output layers is about the same, so information loss is large. However, identity feedforward weights are never observed in the brain; instead, each neuron in the input layer connects to a large number of neurons in the output layer. Using N_x and N_y to denote the number of neurons in the input and output layers, respectively, and K the average number of connections per neuron, the effective noise is reduced by a factor or KN_x^2/N_y^2 (see Eq. 95 below). Thus, if the number of neurons in the output layer is larger than the number in the input layer by a factor much larger than $K^{1/2}$, near complete information transmission is possible. For pure differential correlations, the story is much simpler: so long as the number of neurons in both layers is large, and the added noise doesn't have a strong component in the $\mathbf{f}'(s)$ direction, near complete information transmission always occurs.

4.6.1 Identity feedforward weights

We'll first consider identity feedforward weight, $\mathbf{W} = \mathbf{I}$. We'll start with the matched covariance case. Using Eq. (46), we have

$$\mathbf{\Sigma}_\xi = \frac{I_\eta}{I_x} \mathbf{\Sigma}_\eta. \tag{77}$$

Taking the trace of both sides of this expression gives

$$\frac{I_x}{I_\eta} = \frac{\langle \sigma_\eta^2 \rangle}{\langle \sigma_x^2 \rangle} \tag{78}$$

where $\langle \sigma_x^2 \rangle$ is the average variance of the input layer noise and $\langle \sigma_\eta^2 \rangle$ is the average variance of the added noise. If the added noise is on the same order as the noise in the input layer, information loss is high. Because of synaptic failures and chaotic dynamics, we expect the added noise to be substantial, implying that matching covariances is not an especially good strategy for transmitting information, in the case where $\mathbf{W} = \mathbf{I}$.

Next we consider differential correlations ($\alpha = 0$ in Eq. (46)),

$$\mathbf{\Sigma}_\xi = \frac{I_\eta}{I_x} \frac{\mathbf{f}'\mathbf{f}'}{\mathbf{f}' \cdot \mathbf{\Sigma}_\eta^{-1} \cdot \mathbf{f}'} \tag{79}$$

where we used Eq. (9) for I_η, with $\mathbf{\Sigma}_y$ replaced by $\mathbf{\Sigma}_\eta$. Taking the trace of both sides gives us

$$\frac{I_x}{I_\eta} = \frac{1}{N_x} \frac{\mathbf{f}' \cdot \mathbf{f}'}{\langle \sigma_x^2 \rangle \, \mathbf{f}' \cdot \mathbf{\Sigma}_\eta^{-1} \cdot \mathbf{f}'} \,. \tag{80}$$

If the added noise doesn't have much of a component in the $\mathbf{f}'$ direction, then $\mathbf{f}' \cdot \mathbf{\Sigma}_\eta^{-1} \cdot \mathbf{f}'$ is $\mathcal{O}(N_x)$. In this case, in the large N_x regime, $I_x \ll I_\eta$, and (according to Eq. (76)) information loss is small. In other words, for large neural populations, differential correlations allow small information loss even when the amount of added noise is large.

An especially instructive case is *iid* noise added at the second layer. Using σ_η^2 for its variance, Eq. (80) simplifies to

$$\frac{I_x}{I_\eta} = \frac{1}{N_x} \frac{\sigma_\eta^2}{\langle \sigma_x^2 \rangle} \,. \tag{81}$$

Consequently, for differential correlations and reasonably large neural populations, information loss is relatively small unless the variance in the second layer is *much* larger than the average variance in the first layer (by about a factor of N_x) – something that is not observed in the brain.

Although pure differential correlations can minimize information loss, they are not biologically realistic, as they do not display Poisson-like variability. That's because for differential correlations, the variance of neuron i scales as $f_i'(s)^2$ rather than $f_i(s)$. Fortunately, this can be fixed with very little information loss by adding Poisson-like variability in the input layer. Doing so reduces the information only slightly: for the covariance structure given in Eq. (4a), the information is

$$I_x = \frac{I_0}{1 + \epsilon I_0} \tag{82}$$

where

$$I_0 = \mathbf{f}' \cdot \mathbf{\Sigma}_0^{-1} \cdot \mathbf{f}' \tag{83}$$

is the information associated with the covariance matrix $\mathbf{\Sigma}_0$ (see Sec. 4.7). That information is large whenever $\mathbf{\Sigma}_0$ doesn't contain much of a component in the $\mathbf{f}'$ direction and N_x is large. If these hold, the information in the input layer is approximately equal to $1/\epsilon$ – exactly what it is for pure differential correlations. Moreover, so long as $\mathbf{\Sigma}_\eta$ also doesn't contain much of a component in the $\mathbf{f}'$ direction, information in the output layer is also close to $1/\epsilon$, and very little information is lost. Thus, nearly pure differential correlations are biologically realistic and can lead to very small information loss.

4.6.2 Realistic feedforward weights

For realistic feedforward weights, $\mathbf{W}$, we need to use $\mathbf{\Sigma}_y$ rather than $\mathbf{\Sigma}_\eta$ in Eq. (77), with $\mathbf{\Sigma}_y$ given by Eq. (8). (Note that because the gain function is the identity, $\mathbf{W}_{\text{eff}} = \mathbf{W}$.) We'll start, as above, with the matched covariance case. Taking the trace of both sides of Eq. (77), but with $\mathbf{\Sigma}_\eta$ replaced by $\mathbf{\Sigma}_y$, we have

$$\frac{I_x}{I_\eta} = \frac{\text{tr}[\mathbf{\Sigma}_y]/N_x}{\langle \sigma_x^2 \rangle} \tag{84}$$

where tr denotes trace and, as above, N_x is the number of neurons in the input layer. Using the fact that for any positive semi-definite square $n \times n$ matrix $\mathbf{A}$ (i.e., for any covariance matrix $\mathbf{A}$),

$$\frac{\text{tr}[\mathbf{A}^{-1}]}{n} \geq \frac{n}{\text{tr}[\mathbf{A}]}, \tag{85}$$

we have

$$\frac{I_x}{I_\eta} \geq \frac{1}{\langle \sigma_x^2 \rangle \, \text{tr}[\mathbf{\Sigma}_y^{-1}]/N_x} = \frac{1}{\langle \sigma_x^2 \rangle \, \text{tr}[\mathbf{W}^T \cdot \mathbf{\Sigma}_\eta^{-1} \cdot \mathbf{W}]/N_x} \,, \tag{86}$$

with the second equality following from Eq. (8).

To get a handle on the size of the trace term in the numerator, we note that it can be written

$$\text{tr}[\mathbf{W}^T \cdot \mathbf{\Sigma}_\eta^{-1} \cdot \mathbf{W}] = \text{tr}[\mathbf{W}^T \cdot \mathbf{W}] \left\langle 1/\sigma_\eta^2 \right\rangle_W \tag{87}$$

where, defining $\mathbf{v}_k$ to be the k^{th} eigenvector of $\mathbf{\Sigma}_\eta$, normalized so that $\mathbf{v}_k \cdot \mathbf{v}_k = 1$, and σ_k^2 to be its corresponding eigenvalue,

$$\left\langle 1/\sigma_\eta^2 \right\rangle_W \equiv \frac{1}{\text{tr}[\mathbf{W}^T \cdot \mathbf{W}]} \sum_k \frac{\mathbf{v}_k \cdot \mathbf{W} \cdot \mathbf{W}^T \cdot \mathbf{v}_k}{\sigma_k^2} \,. \tag{88}$$

To see that this really is a weighted average, note that because the $\mathbf{v}_k$ form a complete, orthonormal basis,

$$\sum_k \mathbf{v}_k \cdot \mathbf{W} \cdot \mathbf{W}^T \cdot \mathbf{v}_k = \text{tr}[\mathbf{W} \cdot \mathbf{W}^T]. \tag{89}$$

Inserting Eq. (87) into Eq. (86) gives us

$$\frac{I_x}{I_\eta} \geq \frac{1}{\langle \sigma_\eta^2 \rangle \langle 1/\sigma_\eta^2 \rangle_W} \, \frac{1}{\text{tr}[\mathbf{W}^T \cdot \mathbf{W}]/N_x} \, \frac{\langle \sigma_\eta^2 \rangle}{\langle \sigma_x^2 \rangle} \,. \tag{90}$$

This is similar to Eq. (78), except for two prefactors. The denominator of the first prefactor lies between $\langle \sigma_\eta^2 \rangle / \sigma_{\eta,\max}$ and $\langle \sigma_\eta^2 \rangle / \sigma_{\eta,\min}$. We'll assume this is $\mathcal{O}(1)$ (for *iid* noise it is exactly 1), although we note that it's possible to make it either relatively large or relatively small. The second prefactor is more interesting, as it is the sum of a large number of terms,

$$\frac{\mathrm{tr}[\mathbf{W}^T \cdot \mathbf{W}]}{N_x} = \frac{1}{N_x} \sum_{i=1}^{N_y} \sum_{j=1}^{N_x} W_{ij}^2 \tag{91}$$

where N_y is the number of neurons in the output layer. To determine the size of the weights, we use that fact that

$$\langle y_i \rangle = \sum_{j=1}^{N_x} W_{ij} f_j, \tag{92}$$

and note that $\langle y_i \rangle$ and f_i should be about the same size, on average. Assuming that each neuron in the input layer connects, on average, to K neurons in the output layer, it follows that W_{ij} is nonzero with probability K/N_y. Consequently,

$$\langle y_i \rangle = \sum_j W_{ij} f_j \sim \frac{N_x K}{N_y} W_{\mathrm{typical}} f_{\mathrm{typical}} \tag{93}$$

where W_{typical} and f_{typical} are the typical sizes of the nonzero weights and the f_j, respectively. To ensure that $\langle y_i \rangle$ and f_i are about the same size, we must have

$$W_{\mathrm{typical}} \sim \frac{N_y}{N_x K}\,. \tag{94}$$

Inserting this into Eq. (91), and using the fact that W_{ij} is nonzero with probability K/N_y, we have

$$\frac{\mathrm{tr}[\mathbf{W}^T \cdot \mathbf{W}]}{N_x} \sim \frac{(N_y/N_x)^2}{K} \tag{95}$$

This can be large if $N_y \gg N_x K^{1/2}$. Using this relationship in Eq. (90), we see that information loss can be small in the case of matched covariances, if there is sufficiently large divergence from the input to output layers.

What about differential correlations, $\alpha = 0$? To understand information loss in this case, $\mathbf{\Sigma}_\eta$ is replaced by $\mathbf{\Sigma}_y$ in Eq. (80), giving us

$$\frac{I_x}{I_\eta} = \frac{1}{N_x} \frac{\mathbf{f}' \cdot \mathbf{f}'}{\langle \sigma_x^2 \rangle\, \mathbf{f}' \cdot \mathbf{W}^T \cdot \mathbf{\Sigma}_\eta^{-1} \cdot \mathbf{W} \cdot \mathbf{f}'} \tag{96}$$

where we used Eq. (8) for $\mathbf{\Sigma}_y$. Here the logic is the same as it was in the previous section: so long as $\mathbf{\Sigma}_y$ doesn't have a strong component in the $\mathbf{f}'$ direction, $\mathbf{f}' \cdot \mathbf{W}^T \cdot \mathbf{\Sigma}_\eta^{-1} \cdot \mathbf{W} \cdot \mathbf{f}'$ is $\mathcal{O}(N_y)$, and, since $\mathbf{f}' \cdot \mathbf{f}' \sim \mathcal{O}(N_x)$, information loss is $\mathcal{O}(1/N_y)$. Thus, with realistic feedforward weights, as with the identity case, differential correlations lead to very small information loss in large populations.

4.7 Information in a population with a rank 1 perturbation to the covariance matrix

In the analysis of nonlinear gain functions in Sec. 2.5, it was necessary to construct a covariance matrix such that the information in the first layer was independent of ϵ_u and $\mathbf{u}$. For that we included a prefactor

γ_u in the definition of the covariance matrix, $\mathbf{\Sigma}_\xi$ (see Eq. (12)). Here we determine how γ_u should depend on ϵ_u and $\mathbf{u}$. Our starting point is an expression for the inverse of $\mathbf{\Sigma}_\xi$. As is straightforward to show, via direct substitution, that's given by

$$\mathbf{\Sigma}_\xi^{-1} = \left(\gamma_u\left[\mathbf{\Sigma}_0 + \epsilon_u \mathbf{u}\mathbf{u}\right]\right)^{-1} = \frac{1}{\gamma_u}\left[\mathbf{\Sigma}_0^{-1} - \frac{\epsilon_u \mathbf{\Sigma}_0^{-1}\cdot\mathbf{u}\mathbf{u}\cdot\mathbf{\Sigma}_0}{1+\epsilon_u \mathbf{u}\cdot\mathbf{\Sigma}_0^{-1}\cdot\mathbf{u}}\right]. \tag{97}$$

Thus, the information in the input layer, $\mathbf{f}'\cdot\mathbf{\Sigma}_\xi^{-1}\cdot\mathbf{f}'$, is given by

$$\mathbf{f}'\cdot\mathbf{\Sigma}_\xi^{-1}\cdot\mathbf{f}' = \frac{1}{\gamma_u}\left[\mathbf{f}'\cdot\mathbf{\Sigma}_0^{-1}\cdot\mathbf{f}' - \frac{\epsilon_u(\mathbf{f}'\cdot\mathbf{\Sigma}_0^{-1}\cdot\mathbf{u})^2}{1+\epsilon_u \mathbf{u}\cdot\mathbf{\Sigma}_0^{-1}\cdot\mathbf{u}}\right]. \tag{98}$$

To ensure that this information is independent of γ_u, we let

$$\gamma_u = \frac{1}{I_x}\left[\mathbf{f}'\cdot\mathbf{\Sigma}_0^{-1}\cdot\mathbf{f}' - \frac{\epsilon_u(\mathbf{f}'\cdot\mathbf{\Sigma}_0^{-1}\cdot\mathbf{u})^2}{1+\epsilon_u \mathbf{u}\cdot\mathbf{\Sigma}_0^{-1}\cdot\mathbf{u}}\right]. \tag{99}$$

Note that γ_u depends on s as well as ϵ_u and $\mathbf{u}$.

4.8 Details for Numerical Examples

In this section we provide details for the numerical simulations for each relevant figure.

4.8.1 Figure 2 and its synergistic counterpart, Fig. 7

For the numerical examples in Fig. 2, we generated tuning curves for the first layer of cells using Von Mises distributions [16],

$$f_i(s) = \rho_i + \upsilon_i \exp\left[\beta_i\left(\cos(s-\phi_i)-1\right)\right]. \tag{100}$$

For each cell, the amplitudes, υ, widths, β, peak locations, ϕ, and baseline offsets, ρ, were drawn independently from uniform distributions with the following ranges,

- υ: 1–51
- β: 1–6
- ϕ: 0–2π
- ρ: 0–1

The covariance of the noise in the first layer was given by Eq. (4), with the following parameters,

- blue population: $\epsilon = 10^{-3}$.
- green population: ϵ_u varies with stimulus so that, for each stimulus, the blue and green populations have identical information (on average, $\epsilon_u = 8\times10^{-3}$); $|\mathbf{u}(s)| = |\mathbf{f}'(s)|$; angle between $\mathbf{u}(s)$ and $\mathbf{f}'(s) = 1/8$ of a radian.

With these parameters, the two populations (blue and green) conveyed the same amount of information about the stimulus.

To rule out the possibility that differences in information robustness were due to differences in average correlations within the populations, we forced the average correlations to be the same for the blue and green populations. To do that, we repeatedly took random draws of the parameters describing the tuning curves (ρ, v, β and ϕ) until the population averaged correlations matched between the two populations. This resulted in average correlations of -7×10^{-5}, and we used this set of tuning curves for our subsequent information calculations.

We computed the information, $I_y(s)$, in the second-layer responses using Eq. (3b), with $g(x) = x$, $\mathbf{W} = \mathbf{I}$, and $\mathbf{\Sigma}_\eta = \sigma^2\mathbf{I}$. For the trial-shuffled information (Fig. 2C), we used Eq. (3a), with all off-diagonal elements of the covariance matrices $\mathbf{\Sigma}_\xi$ set to zero. For all of these information calculations, we computed the information, $I_x(s)$ or $I_y(s)$, for 100 different stimulus values s, uniformly spaced between 0 and 2π, and then averaged over these 100 different values.

To assess whether synergistic population codes can similarly vary in their robustness to corruption by noise, we repeated our calculations from Fig. 2, but modified the covariance matrices to make the population synergistic (Fig. 7C: the correlated responses convey more stimulus information than would independent cells with the same variances). To do that we again used the covariance matrices given in Eq. (4), but we made ϵ and ϵ_u negative: $\epsilon = -5 \times 10^{-4}$ and $\langle\epsilon_u\rangle = -3 \times 10^{-4}$ (as in Fig. 2, ϵ_u depends on the stimulus, s: it was chosen so that for each value of s the blue and green populations have identical stimulus information). We chose $\mathbf{u}(s)$ so that it had the same magnitude as $\mathbf{f}'(s)$ and made an angle of 1/4 of a radian with $\mathbf{f}'(s)$. We used the same functions and distributions for the tuning curves as in Fig. 2, but used a different seed for the random number generator. As in Fig. 2, the seed was chosen (via multiple draws of the tuning curve parameters) so that the two populations had the same average correlations (in this case 2×10^{-5}). Also as in Fig. 2, the populations were roughly Poisson-like, in the sense that the mean and variance of the activity of each neuron was approximately equal. (Both the "green" and the "blue" populations have average Fano factors – averaged over neurons and stimuli – of 0.99.) We again found that equally-informative population codes could vary significantly in terms of their robustness to noise (Fig. 7B).

4.8.2 Figure 5

To generate Fig. 5B, we analytically computed the means of the second layer responses, resulting in the expression

$$\mu_i(s) = \Phi\left[\frac{f_i(s) - \theta_i}{\sigma_i(s)}\right], \tag{101}$$

where θ_i is the i^{th} cell's firing threshold, σ_i is the standard deviation of the input noise to the cell, and $\Phi(\cdot)$ is the Gaussian cumulative distribution function. For each cell, the input function $f_i(s)$ was given by a Von Mises distribution, Eq. (100) (with the same distribution of parameters – v, β, ϕ and ρ – as in the preceding examples), and the spiking threshold, θ_i, was set to 3/4 of the peak height of the input tuning curve: $\theta_i = 3\left(\rho_i + \upsilon_i\right)/4$.

It is not straightforward to compute the covariance matrix of correlated responses generated by the dichotomized Gaussian model, so we used Monte Carlo methods to estimate the covariance: we took 10^6 draws from the distribution of $\mathbf{x}$, and for each draw we computed the corresponding responses, $\mathbf{y}$, using the thresholding operation (Eq. (11)). We then computed the covariance of these simulated responses, and used them to estimate the linear Fisher information in the second layer activities via the standard expression,

$$I_y(s) = \frac{\partial\boldsymbol{\mu}(s)}{\partial s} \cdot \mathrm{Cov}(\mathbf{y}|s)^{-1} \cdot \frac{\partial\boldsymbol{\mu}(s)}{\partial s}. \tag{102}$$

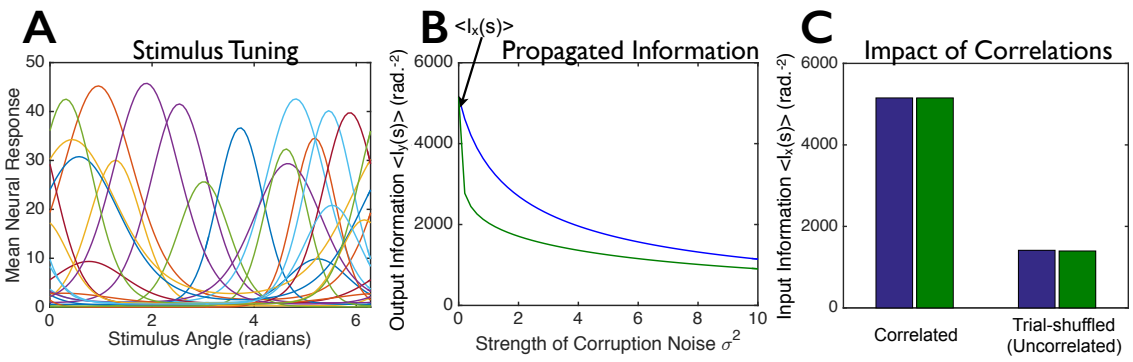

Figure 7. Not all synergistic population codes are equally robust against corruption by noise. This figure is similar to Fig. 2, but with synergistic instead of redundant population codes. We constructed two model populations – each with the same 100 tuning curves (20 randomly-chosen example tuning curves are shown in panel **A**) – for the first layer of cells. The two populations have different covariance structures $\mathbf{\Sigma}_\xi$ for their trial-to-trial variability (see main text, Eq. (4)), but convey identical amounts of information, $I_x(s)$, about the stimulus. (**B**) We corrupted the responses of each neural population by Gaussian noise (independently and identically distributed for all cells) of variance σ^2, to mimic corruption that might arise as the signals propagate through a multi-layered neural circuit, and computed the output information, $I_y(s)$, that these further-corrupted responses convey about the stimulus (blue and green curves). (**C**) Input information $I_x(s)$ in the two model populations (left; "correlated") and information that would be conveyed by the model populations if they had their same tuning curves and levels of trial-to-trial variability, but no correlations between cells (right; "trial-shuffled"). For panels B and C, we computed the information for 100 different stimulus values, equally spaced between 0 and 2π, and averaged the information over these stimuli.

4.8.3 Figure 6

Figure 6 was made in the same fashion as Fig. 5, with the exception that noise was added before the spike generation nonlinearity. The noise, $\boldsymbol{\zeta}$, was Gaussian and drawn *iid* with variance σ_ζ^2.

Acknowledgments

We thank Robert Townley, Kresimir Josic, Fred Rieke, Braden Brinkman, Maxwell Turner, and Alison Weber for helpful comments on the project. JZ's contribution to this work was partially supported by an Azrieli Global Scholar Award from the Canadian Institute For Advanced Research (CIFAR). PEL was supported by the Gatsby Charitable Foundation. ESB acknowledges the support of NSF Grant CRCNS-1208027 and a Simons Fellowship in Mathematics, and thanks the Allen Institute founders, Paul G. Allen and Jody Allen, for their vision, encouragement and support.

References

1. Britten K, Shadlen M, Newsome W, Movshon J (1993) Responses of neurons in macaque MT to stochastic motion signals. Visual Neurosci 10: 1157-1169.

2. Softky W, Koch C (1993) The highly irregular firing of cortical cells is inconsistent with temporal integration of random epsp's. J Neurosci 13: 334–350.

3. Faisal A, Selen L, Wolpert D (2008) Noise in the nervous system. Nat Rev Neurosci 9: 292-303.

4. Churchland M, Yu B, Cunningham J, Sugrue L, Cohen M, et al. (2010) Stimulus onset quenches neural variability: a widespread cortical phenomenon. Nat Neurosci 13: 369-378.

5. Franke F, Fiscella M, Sevelev M, Roska B, Hierlemann A, et al. (2016) Structure of neural correlation and how they favor coding. Neuron 89: 409-422.

6. Zylberberg J, Cafaro J, Turner M, Shea-Brown E, Rieke F (2016) Direction-selective circuits shape noise to ensure a precise population code. Neuron (89): 369-383.

7. Zylberberg J, Hyde R, Strowbridge B (2016) Dynamics of robust pattern separability in the hippocampal dentate gyrus. Hippocampus (29): 623-632.

8. Zohary E, Shadlen MN, Newsome WT (1994) Correlated neuronal discharge rate and its implications for psychophysical performance. Nature 370: 140–143.

9. Abbott LF, Dayan P (1999) The effect of correlated variability on the accuracy of a population code. Neural Comput 11: 91–101.

10. Sompolinsky H, Yoon H, Kang K, Shamir M (2001) Population coding in neuronal systems with correlated noise. Phys Rev E 64: 051904.

11. Romo R, Hernandez A, Zainos A, Salinas E (2003) Correlated neuronal discharges that increase coding efficiency during perceptual discrimination. Neuron 38: 649–657.

12. Averbeck BB, Latham PE, Pouget A (2006) Neural correlations, population coding and computation. Nat Rev Neurosci 7: 358–366.

13. Shamir M, Sompolinsky H (2006) Implications of neuronal diversity on population coding. Neural Comput 18: 1951–1986.

14. Averbeck BB, Lee D (2006) Effects of noise correlations on information encoding and decoding. J Neurophys 95: 3633–3644.

15. Josić K, Shea-Brown E, Doiron B, de la Rocha J (2009) Stimulus-dependent correlations and population codes. Neural Comput 21: 2774–2804.

16. Ecker AS, Berens P, Tolias AS, Bethge M (2011) The Effect of Noise Correlations in Populations of Diversely Tuned Neurons. J Neurosci 31: 14272–14283.

17. Cohen MR, Kohn A (2011) Measuring and interpreting neuronal correlations. Nat Neurosci 14: 811–819.

18. Latham P, Roudi Y (2011) Role of correlations in population coding. arXiv:11096524 [q-bio/NC] .

19. da Silveira RA, Berry MJ (2014) High-Fidelity Coding with Correlated Neurons. PLoS Comput Biol 10: e1003970.

20. Hu Y, Zylberberg J, Shea-Brown E (2014) The sign rule and beyond: Boundary effects, flexibility, and noise correlations in neural population codes. PLoS Comput Biol 10: e1003469.

21. Shamir M (2014) Emerging principles of population coding: in search for the neural code. Curr Opin Neurobiol 25: 140-148.

22. Moreno-Bote R, Beck J, Kanitscheider I, Pitkow X, Latham P, et al. (2014) Information-limiting correlations. Nat Neurosci 17: 1410-1417.

23. Zylberberg J, Shea-Brown E (2015) Input nonlinearities can shape beyond-pairwise correlations and improve information transmission by neural populations. Phys Rev E 92: 062707.

24. Cayco-Gajic A, Zylberberg J, Shea-Brown E (2015) Triplet correlations among similarly-tuned cells impact population coding. Front Comput Neurosci 9: 57.

25. Pouget A, Deneve S, Ducom JC, Latham P (1999) Narrow versus wide tuning curves: what's best for a population code? Neural Comput 11: 85-90.

26. Zhang K, Sejnowski TJ (1999) Neuronal tuning: To sharpen or broaden? Neural Comput 11: 75–84.

27. Wilke SD, Eurich CW (2002) Representational accuracy of stochastic neural populations. Neural Comput 14: 155–189.

28. Tkačik G, Prentice J, Balasubramanian V, Schneidman E (2010) Optimal population coding by noisy spiking neurons. Proc Natl Acad Sci USA 107: 14419–14424.

29. Beck J, Bejjanki VR, Pouget A (2011) Insights from a simple expression for linear fisher information in a recurrently connected population of spiking neurons. Neural Comput 23: 1484–1502.

30. Toyoizumi T, Aihara K, Amari SI (2006) Fisher information for spike-based population decoding. Phys Rev Lett 97: 098102.

31. Bejjanki V, Beck J, Lu Z, Pouget A (2011) Perceptual learning as improved probabilistic inference. Nat Neurosci 14: 642–648.

32. Salinas E, Sejnowski TJ (2000) Impact of correlated synaptic input on output firing rate and variability in simple neuronal models. J Neurosci 20: 6193–6209.

33. Salinas E, Sejnowski T (2001) Correlated neuronal activity and the flow of neural information. Nat Rev Neurosci 2: 539-550.

34. Reid R (2001) Divergence and reconvergence: multielectrode analysis of feedforward connections in the visual system. Prog Brain Res 130: 141-154.

35. Bruno R (2011) Synchrony in sensation. Curr Opin Neurobiol 21: 701-708.

36. Abeles M (1982) Role of the cortical neuron: integrator or coincidence detector? Isr J Med Sci 18: 83-92.

37. Seriès P, Latham PE, Pouget A (2004) Tuning curve sharpening for orientation selectivity: coding efficiency and the impact of correlations. Nat Neurosci 7: 1129–1135.

38. Renart A, van Rossum M (2011) Transmission of population-coded information. Neural Comput 24: 391-407.

39. Lampl I, Reichova I, Ferster D (1999) Synchronous membrane potential fluctuations in neurons of the cat visual cortex. Neuron 22: 361-374.

40. Alonso J, Usrey W, Reid R (1996) Precisely correlated firing of cells in the lateral geniculate nucleus. Nature 383: 815–819.

41. Goris R, Movshon J, Simoncelli E (2014) Partitioning neuronal variability. Nat Neurosci 17: 858-865.

42. Smith M, Kohn A (2008) Spatial and temporal scales of neuronal correlation in primary visual cortex. J Neurosci 28: 12591-12603.

43. Ecker A, et al (2014) State dependence of noise correlations in macaque primary visual cortex. Neuron 82: 235-248.

44. Scholvinck M, Saleem A, Benucci A, Harris K, Carandini M (2015) Cortical state determines global variability and correlations in visual cortex. J Neurosci 35: 170-178.

45. Lin IC, Okun M, Carandini M, Harris K (2015) The nature of shared cortical variability. Neuron 87: 644-656.

46. Hornik K, Stinchcombe M, White H (1989) Multilayer feedforward networks are universal function approximators. Neural Netw 2: 359–366.

47. Barlow H, Levick W (1965) The mechanism of directionally selective units in rabbit's retina. J Physiol 178: 477-504.

48. Cramer H (1946) Mathematical methods of statistics. Princeton University Press.

49. Rao C (1945) Information and the accuracy attainable in the estimation of statistical parameters. Bull Calcutta Math Soc 37: 81-89.

50. Schneidman E, Still S, Berry MJ, Bialek W (2003) Network Information and Connected Correlations. Physical Review Letters 91: 238701.

51. Adibi M, McDonald J, Clifford C, Arabzadeh E (2013) Adaptation improves neural coding efficiency despite increasing correlations in variability. J Neurosci 33: 2108-2120.

52. Macke JH, Berens P, Ecker AS, Tolias AS, Bethge M (2009) Generating spike trains with specified correlation coefficients. Neural Comput 21: 397–423.

53. Macke J, Opper M, Bethge M (2011) Common Input Explains Higher-Order Correlations and Entropy in a Simple Model of Neural Population Activity. Phys Rev Lett 106: 208102.

54. Yu S, Yang H, Nakahara H, Santos G, Nikolic D, et al. (2011) Higher-order correlations characterized in cortical activity. J Neurosci 31: 17514-17526.

55. Amari SI, Nakahara H, Wu S, Sakai Y (2003) Synchronous firing and higher-order interactions in neuron pool. Neural Comput 15: 127–142.

56. Bethge M, Berens P (2008) Near-maximum entropy models for binary neural representations of natural images, JC Platt, D. Koller, Y. Singer, S. Roweis, Editors. Advances in neural information processing systems 20.

57. Leen D, Shea-Brown E (2015) A simple mechanism for beyond-pairwise correlations in integrate-and-fire neurons. J Math Neurosci 5: 17.

58. Barlow H (2001) Redundancy reduction revisited. Network-Computation in Neural Systems 12: 241–253.

59. Beck JM, Ma W, Pitkow X, Latham PE, Pouget A (2012) Not noisy, just wrong: the role of suboptimal inference in behavioral variability. Neuron 74: 30–39.

60. Kanitscheider I, Coen-Cagli R, Pouget A (2015) Origin of information-limiting noise correlations. Proc Natl Acad Sci USA 112: E6973-E6982.

61. Graf A, Kohn A, Jazayeri M, Movshon J (2011) Decoding the activity of neuronal populations in macaque primary visual cortex. Nat Neurosci 14: 239–245.

62. Berens P, Ecker A, Cotton R, Ma W, Bethge M, et al. (2012) A fast and simple population code for orientation in primate v1. J Neurosci 32: 10618–10626.